Universidade do Minho
Escola de Economia e Gestão

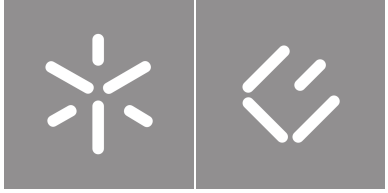
Cátia Cristina Afonso Cerqueira

Technology adoption, labour market institutions, and firm capabilities

March, 2026

Technology adoption, labour market institutions, and firm capabilities
Cátia Cerqueira
UMinho | 2026

REPÚBLICA PORTUGUESA
CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR
fct Fundação para a Ciência e a Tecnologia



Universidade do Minho
Escola de Economia e Gestão

Cátia Cristina Afonso Cerqueira

Technology adoption, labour market institutions, and firm capabilities

Doctorate Thesis
Doctorate in Economics

Work developed under the supervision of:
Professor Miguel Portela
Professor Fernando Alexandre

March, 2026

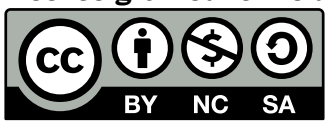
COPYRIGHT AND TERMS OF USE OF THIS WORK BY A THIRD PARTY

This is academic work that can be used by third parties as long as internationally accepted rules and good practices regarding copyright and related rights are respected.

Accordingly, this work may be used under the license provided below.

If the user needs permission to make use of the work under conditions not provided for in the indicated licensing, they should contact the author through the RepositoriUM of Universidade do Minho.

License granted to the users of this work



Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International
CC BY-NC-SA 4.0

https://creativecommons.org/licenses/by-nc-sa/4.0/deed.en

This document was created with the (pdf/Xe/Lua)LaTeX processor and the NOVAthesis template (v7.1.7) (Lourenço, 2021).

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Professor Miguel Portela, for his exceptional guidance, continuous support, and meticulous review throughout the course of this doctoral journey. His insightful feedback, encouragement, and dedication were fundamental to the completion of this thesis and ensured its final quality.

I am also sincerely grateful to Professor Fernando Alexandre, whose guidance was instrumental in the development of the research presented in this work. I am particularly thankful for his support and encouragement during the early stages, as well as for his contributions to the formulation of key ideas, which were invaluable in shaping the direction of this project.

I would like to extend a special thanks to my PhD colleague Hélder Costa for the many conversations, exchanges of ideas, and shared knowledge along the way. His willingness to share information and his encouragement during more challenging moments were instrumental in helping me bring this work to completion.

My deepest gratitude goes to my husband, Pedro Costa, and to my daughter, Clara Costa, who was born during the course of this journey. Their love, patience, and understanding in the face of my many absences gave me the strength and motivation to persevere.

I am also profoundly grateful to my parents and my sisters for their unconditional support.

I gratefully acknowledge the support of the research centre NIPE (Centre for Research in Economics and Management) and the financial support provided by Fundação para a Ciência e a Tecnologia (FCT) through the PhD scholarship UI/BD/150752/2020.

Finally, I would like to acknowledge the members of the academic community at the University of Minho who, directly or indirectly, contributed to the development of this work.

This thesis is the result not only of individual effort, but also of the support, encouragement, and generosity of many people to whom I am deeply grateful.

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho.

DECLARATION ON THE USE OF ARTIFICIAL INTELLIGENCE (AI) TOOLS

I hereby declare having conducted this academic work with integrity and confirm that I have not resorted to any form of undue use of AI tools, nor to the falsification or omission of information, sources, data, results or processes in any of the stages leading to its elaboration.

I further declare that any use of AI tools was done in a transparent and responsible manner, being duly noted in the work, and that I critically verified the accuracy of the content produced with their support, assuming full responsibility for the submitted content.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho, as well as the applicable standards of academic integrity and data protection.

Resumo

Adoção de tecnologia, instituições do mercado de trabalho e capacidades empresariais

Esta tese analisa como as capacidades das empresas e as instituições do mercado de trabalho moldam a adoção de tecnologias digitais e de automação, bem como os seus efeitos no desempenho das empresas e nos resultados dos trabalhadores. Compreender os determinantes e as consequências da adoção tecnológica é essencial para analisar a produtividade, a transformação estrutural e a distribuição de riqueza. A análise utiliza uma base de dados detalhada que liga empresas e trabalhadores em Portugal, combinando o inquérito à utilização de tecnologias de informação e comunicação nas empresas (IUTICE), registos de comércio internacional (CI), sistema de contas integradas das empresas (SCIE) e quadros de pessoal (QP).

O primeiro capítulo estuda o momento e as consequências da adoção de tecnologias de computação em nuvem e de *big data*, distinguindo entre pioneiros, seguidores e empresas que não adotam. Empresas com gestores e trabalhadores mais qualificados, bem como gestores com menor antiguidade, tendem a ser pioneiros na adoção destas tecnologias, especialmente no caso de *big data*. A adoção está associada a um aumento do desempenho das empresas. Para *big data*, o aumento do desempenho é ampliado por características do capital humano, em particular a educação, especialmente dos gestores, e os pioneiros mantêm vantagens persistentes em produtividade e valor acrescentado em relação aos seguidores. Os resultados sugerem que o papel do capital humano na mediação do impacto da tecnologia no desempenho, assim como as potenciais vantagens de *timing*, dependem da complexidade da tecnologia adotada.

O segundo capítulo analisa o impacto da reforma da escolaridade mínima obrigatória de 2009, que aumentou a idade mínima de saída do ensino de 15 para 18 anos (ou a conclusão do 12º ano de escolaridade) e restringiu o acesso das empresas a trabalhadores menores de idade, avaliando se esta escassez de mão de obra induziu a automação no setor da manufatura. Através da construção de uma base de dados rica que liga características de empresas, trabalhadores e dados de comércio internacional, comparamos as alterações no período pós-reforma entre empresas do setor da manufatura com maior e menor exposição prévia ao trabalho de menores. As empresas mais expostas registaram uma redução superior do emprego de menores e ajustaram-se sobretudo ao nível da composição da força de trabalho,

com aumentos no emprego em ocupações associadas à automação e tecnologias de informação e comunicação (TIC). Em contraste, não existe evidência robusta de aumento das importações de robôs. No entanto, o aumento da probabilidade de contratação de trabalhadores de TIC pode não ser específico da reforma, uma vez que exercícios de *placebo* também produzem estimativas positivas e estatisticamente significativas. No conjunto, os resultados apontam para um ajustamento das empresas mais expostas após a reforma, concentrado na composição da força de trabalho.

O terceiro capítulo investiga se os aumentos do salário mínimo influenciaram a adoção de robôs, considerando a alteração do preço do fator de produção trabalho, bem como os impactos da automação sobre empresas e trabalhadores. Empresas mais expostas a aumentos do salário mínimo apresentam maior probabilidade de adquirir robôs, embora com um desfasamento de três a cinco anos após o choque. A adoção está associada a uma expansão da escala das empresas, incluindo aumentos de produtividade e valor acrescentado, embora parcialmente determinados por efeitos de auto-seleção, o que motivou a estratégia empírica centrada nos impactos sobre trabalhadores dentro da mesma empresa. Trabalhadores em tarefas mais expostas à automação com contratos a termo têm maior probabilidade de sair da empresa, após a adoção de robôs, enquanto que trabalhadores sem termo são realocados internamente para ocupações menos expostas. Paralelamente, trabalhadores em ocupações mais expostas têm uma penalização salarial decorrente da adoção de robôs pela empresa empregadora.

De forma geral, a tese demonstra que a adoção tecnológica depende da interação entre as capacidades internas das empresas e as instituições do mercado de trabalho. A qualidade da gestão e as competências da força de trabalho determinam a capacidade de implementar tecnologias digitais complexas, enquanto que alterações institucionais na oferta e nos custos laborais moldam os incentivos à automação e os seus impactos sobre empresas e trabalhadores. Do ponto de vista da política pública, educação, regulação salarial e inovação devem ser consideradas de forma integrada, através de políticas complementares que promovam competências, apoiem a mobilidade laboral e assegurem que os ganhos de produtividade decorrentes da digitalização e da automação sejam amplamente partilhados.

Palavras-chave: Automação, Digitalização, Instituições do Mercado de Trabalho, Produtividade, Trabalhadores

Abstract

Technology adoption, labour market institutions, and firm capabilities

This thesis examines how firm capabilities and labour market institutions shape the adoption of digital and automation technologies, and how these adoption decisions affect firm performance and worker outcomes. Understanding both the determinants and consequences of technology adoption is essential for analysing productivity dynamics, structural transformation, and distributional effects. The analysis draws on a linked employer–employee database of Portuguese firms, combining the survey on the use of information and communication technologies in enterprises (IUTICE), international trade records (CI), the integrated accounts system of enterprises (SCIE), and personnel records (QP).

The first chapter investigates the timing and consequences of adopting cloud computing and big data technologies, distinguishing between first movers, followers, and non–adopters. Firms with more educated managers and workers, and managers with shorter tenure, adopt earlier, particularly in the case of big data. Adoption is associated with improved firm performance, and the benefits of big data are amplified by stronger human capital. For big data, early adopters retain persistent productivity and value-added advantages over followers, highlighting that technological complexity determines the role of timing and internal capabilities in shaping performance effects.

The second chapter analyses the impact of Portugal's 2009 compulsory schooling reform on firms' adjustment towards automation, which raised the school–leaving age from 15 to 18 (or to the completion of 12 years of education) and restricted firms' access to underage workers. Using linked employer–employee and trade data, we compare post-reform changes between manufacturing firms with greater and lower pre-reform exposure to underage labour. More exposed firms experienced a decline in underage employment and adjusted primarily through changes in workforce composition, with increases in complementary ICT employment and some evidence of occupational reallocation towards automation-related tasks. However, the increase in the probability of hiring ICT workers does not appear to be reform-specific, as placebo tests also produce positive and statistically significant estimates. Overall, the findings suggest heterogeneous firm responses to a labour supply shock, primarily concentrated in labour reallocation and occupational restructuring rather than in capital deepening or automation via robot imports.

The third chapter examines whether statutory minimum wage increases influence robot adoption and

the consequent effects of automation on firm and worker outcomes. Firms exposed to larger wage increases are more likely to adopt industrial robots, with effects materialising with a lag of three to five years. Robot adoption is associated with expansion in firm scale, including value added and employment, although these trajectories are partly driven by positive selection, as high-performing firms are more likely to adopt. These firm-level patterns motivate our empirical strategy, which focuses on differences in outcomes among workers within the same firm. Workers on fixed-term contracts in tasks more exposed to automation face higher separation rates, while permanent workers are reassigned internally. Workers in more automatable tasks employed in robot-adopting firms experience a relative wage penalty compared with less-exposed workers.

Overall, the thesis shows that technology adoption is shaped by the interaction between firm capabilities and labour market institutions. Managerial quality and workforce skills determine firms' capacity to adopt complex digital technologies, while institutional changes in labour supply and costs influence incentives for automation and workers' labour market outcomes. From a policy perspective, education, wage, and innovation policies should be considered jointly, with complementary measures to promote skill formation, support worker mobility, and ensure productivity gains from digitalisation and automation are broadly shared.

Keywords: Automation, Digitalisation, Labour Market Institutions, Productivity, Workers

Contents

List of Figures xii

List of Tables xiv

Acronyms xvii

1 Introduction 1

2 Timing of cloud computing and big data adoption: Cohort characteristics and firm performance 5

2.1 Introduction . . . 5
2.2 Literature review . . . 7
2.2.1 Technology adoption . . . 7
2.2.2 Impacts on firm performance and aggregate productivity . . . 8
2.2.3 Role of management . . . 9
2.2.4 Significance of timing . . . 10
2.3 Conceptual framework and hypotheses . . . 10
2.4 Data and variables description . . . 12
2.4.1 Data . . . 12
2.4.2 Variable description . . . 14
2.4.3 Descriptive statistics . . . 15
2.5 Methodology and econometric strategy . . . 19
2.5.1 Characterizing cloud computing and big data adopters . . . 19
2.5.2 The impact of cloud computing and big data on firm performance . . . 20
2.6 Results . . . 23
2.6.1 Characterization of digital adopters . . . 23
2.6.2 The impact of digitalization on firm performance . . . 26
2.7 Concluding remarks . . . 33

3 Keeping teenagers in school and firm adjustment: evidence from Portuguese manufacturing firms **35**
3.1 Introduction . . . 35
3.2 Related literature . . . 38
3.2.1 Compulsory schooling reforms and labour supply . . . 38
3.2.2 Automation, robots, and firm behaviour . . . 39
3.2.3 Labour scarcity and induced technological change . . . 40
3.3 Institutional context . . . 41
3.3.1 The 2009 reform in context . . . 41
3.3.2 Enrolment response . . . 42
3.3.3 Young workers in manufacturing . . . 43
3.4 Data and descriptive statistics . . . 44
3.4.1 Data sources . . . 44
3.4.2 Descriptive statistics . . . 47
3.5 Research design and empirical strategy . . . 50
3.5.1 Identification strategy . . . 50
3.5.2 Firm-level specifications . . . 52
3.5.3 Outcomes, exposure, and sample construction . . . 54
3.5.4 Heterogeneity and mechanisms . . . 56
3.5.5 Supporting diagnostics and supplementary estimators . . . 57
3.5.6 Attrition and aggregate checks . . . 58
3.5.7 Threats to identification . . . 58
3.6 Results . . . 59
3.6.1 First stage: the reform reduced underage employment . . . 60
3.6.2 Event-study evidence . . . 61
3.6.3 Firm-level adjustment margins: robot imports, occupational reorganization, and Information and Communication Technologies (ICT) employment . . . 62
3.6.4 Heterogeneity by firm size and labour intensity . . . 66
3.6.5 Firm exit and attrition robustness . . . 68
3.7 Conclusion . . . 69

4 Minimum wage, robot adoption, and worker adjustment: Evidence from Portugal **75**
4.1 Introduction . . . 75
4.2 Literature review . . . 77
4.2.1 Minimum wage, labour costs, and robot adoption . . . 77
4.2.2 Automation and firm-level outcomes . . . 79
4.2.3 Worker-level effects and task reallocation . . . 80
4.3 Data, variables description, and summary statistics . . . 81

4.3.1 Data . . . 82
4.3.2 Sample and variables description . . . 82
4.3.3 Summary statistics . . . 83
4.4 Empirical strategy . . . 86
4.4.1 Step 1: Minimum wage and robot adoption . . . 87
4.4.2 Step 2: Robot adoption and worker-level outcomes . . . 88
4.5 Results . . . 90
4.5.1 Step 1: Minimum wage and robot adoption . . . 91
4.5.2 Step 2: Worker level outcomes . . . 93
4.6 Conclusion . . . 105

5 Conclusion **107**

Bibliography **111**

Appendices

A Timing of cloud computing and big data adoption: Cohort characteristics and firm performance **121**
A.1 Data and Variable Description . . . 121
A.2 Results Without Propensity Score Reweighting . . . 127

B Keeping teenagers in school and firm adjustment: evidence from Portuguese manufacturing firms **132**
B.1 Additional descriptive statistics . . . 132
B.2 Robustness checks . . . 132
B.3 Placebo exercises . . . 155

C Minimum wage, robot adoption, and worker adjustment: Evidence from Portugal **162**
C.1 Additional descriptive statistics . . . 162
C.2 Additional results – minimum wage and robot adoption . . . 165
C.3 Descriptive firm-level trajectories around robot adoption . . . 170
C.4 Additional results – worker level outcomes . . . 177

List of Figures

1 Enrolment rates relative to target-age population, 1970–2023 (%) 36
2 High-school enrolment rate relative to resident population aged 15–18 by Nomenclature of territorial units for statistics – level 2 (NUTS 2) region, 2000–2023 (%) 43
3 Education composition of new manufacturing entrants, 1991–2021 45
4 Mean annual robot, automation, occupation-proxy, and ICT outcomes for more exposed and less exposed firms, 1996–2021 (Tier 1, Tier 2, Tier 3, occupation proxies, and ICT; exposure defined as any worker under age 18 during 2006–2008) 60
5 Event study for annual import, occupation, and ICT outcomes under the baseline binary exposure (Panel A) 63

6 Cumulative robot stock in manufacturing and minimum wage levels in Portugal, 2010–2021 84
7 Minimum wage and robot adoption – local projection coefficients 92
8 Minimum wage and robot adoption – differences in differences estimation 94
9 Robot adoption, exposure, and worker wages – event study 96
10 Robot adoption, exposure, and worker displacement – fixed term workers, event study . 100
11 Robot adoption, exposure, and worker reallocation – all workers, event study 101

12 Flowchart illustrating the database construction process 122

A1 Strict robot import dynamics (Tier 1), 1996–2021 135
B1 Event study: cumulative import outcomes across Tier 1, Tier 2, and Tier 3 (Panel A) . . 151
B2 Mean cumulative automation outcomes for more and less exposed firms, 1996–2021 (Tier 1, Tier 2, Tier 3; exposure defined as any worker under age 18 during 2006–2008) 152
B3 Annual outcomes under alternative exposure definitions (binary, 1%, and 10% cutoffs) . . 153
B4 Cumulative outcomes under alternative exposure definitions (binary, 1%, and 10% cutoffs) 154

C1 Minimum wage and robot adoption , purchase in t – local projection coefficients 165
C2 Robot adoption and productivity – descriptive event-study trajectories 172
C3 Robot adoption and value added – descriptive event-study trajectories 173

C4 Robot adoption and employment – descriptive event-study trajectories 174
C5 Robot adoption and labour share – descriptive event-study trajectories 175
C6 Robot adoption and exports – descriptive event-study trajectories 176
C7 Robot adoption, exposure, and worker reallocation – fixed-term workers, event study . . 177

List of Tables

2 Percentage of firms that have adopted technology, broken down by size and year (%) . . 16
3 Descriptive statistics – Cloud Computing 18
4 Descriptive statistics – Big Data 19
5 Selection into digitalization adoption – Multinomial logit analysis 24
6 First movers versus non-adopters – Impact of Cloud Computing 27
7 First movers versus followers – Long-term impact of Cloud Computing 28
8 First movers versus non-adopters – Impact of Big Data 31
9 First movers versus followers – Long-term impact of Big Data 32

10 Mandatory schooling reforms in Portugal 36
11 Regional post-reform shifts in high-school enrolment (NUTS 2) 44
12 Variable descriptions 48
13 Variable descriptions (continued) 49
14 Descriptive statistics aligned with the empirical design 51
15 First-stage firm-level Difference-in-Differences (DiD) estimates for minor employment . . 61
16 Firm-level DiD estimates for four annual adjustment outcomes 67
17 Heterogeneity by firm size and labour intensity in four annual adjustment outcomes . . . 71
18 Heterogeneity by firm size and labour intensity in four annual adjustment outcomes (continued) 72
19 Firm-level DiD estimates for firm exit in the Integrated Accounts System of Enterprises (SCIE)-linked sample 73
20 Balanced-panel robustness: main adjustment outcomes under survivor restrictions . . . 74

21 Variable descriptions and sources 85
22 Summary statistics on adopters and non-adopters 86
23 Minimum wage and robot adoption 93
24 Robot adoption, exposure, and worker wages 95
25 Robot adoption, exposure, and worker displacement – fixed term workers 97

26 Robot adoption, exposure, and worker displacement – all workers 98
27 Robot adoption, exposure, and worker displacement – open-ended contract workers . . 99
28 Robot adoption, exposure, and worker reallocation – fixed term 102
29 Robot adoption, exposure, and worker reallocation – all workers 103
30 Robot adoption, exposure, and worker reallocation – open-ended contract workers . . . 104

31 Variable description 123
32 Share of technological adopters by economic sector (%) 124
33 Distribution of technological adopters by sector (%) 125
34 Distribution of technological cloud computing by main economic sector (%) 126
35 Distribution of technological big data by main sector (%) 126
36 First movers versus non-adopters – Impact of Cloud Computing (without propensity score reweighting) 128
37 First movers versus followers – Long-term impact of Cloud Computing (without propensity score reweighting) 129
38 First movers versus non-adopters – Impact of Big Data (without propensity score reweighting 130
39 First movers versus followers – Long-term impact of Big Data (without propensity score reweighting) 131

A1 Descriptive statistics for firms established before and after the 2009 reform 133
A2 Strict Tier 1 robot adoption rates by firm size, 1996–2021 134
A3 Strict Tier 1 robot adoption rates by manufacturing sector 134
A4 Distribution of strict Tier 1 robot adopters across manufacturing sectors 135
B1 Pooled firm-level DiD estimates without firm fixed effects 137
B2 Pooled firm-level DiD estimates without firm fixed effects (continued) 138
B3 Pooled logit estimates for firm exit without firm fixed effects 139
B4 Five-outcome firm-level DiD estimates with annual and cumulative import outcomes . . 140
B5 Cumulative import outcomes under the inverse hyperbolic sine transformation 141
B6 Four annual outcomes under the original-paper control set 142
B7 Five annual and cumulative outcomes under the original-paper control set 143
B8 Binary-exposure estimates under alternative baseline control sets 144
B9 DiD estimates without sector-year, region-year, or baseline interacted controls 145
B10 Probit coefficients for annual robot-importer and Two-Way Fixed Effects (TWFE)-worker outcomes 146
B11 Firm-level DiD estimates on the 2005–2015 window 147
B12 Poisson Pseudo-Maximum Likelihood (PPML) sensitivity for annual robot importing and cumulative robot-import value 148
B13 Annual comparison across import tiers and occupation proxies 148

B14 Cumulative comparison across import tiers 149
B15 Supplementary event-study lead-test p-values for the four main annual adjustment outcomes 149
B16 Threshold sensitivity for binary exposure in four annual adjustment outcomes 150
B17 Binary-exposure estimates under imputation and trend-augmented specifications 155
B18 Matched common-support balance diagnostics 156
B19 Matched common-support estimates for binary exposure 157
B20 Standard-error comparison under alternative clustering 158
B21 Sector–region cell-level DiD estimates 159
C1 Placebo estimates using negative-control exposure variables 160
C2 Placebo timing estimates with fake reform years in the pre-2009 window 161

C3 Descriptive statistics adopters and non-adopters by size 162
C4 Within-firm variation in occupational exposure measures (2010) 164
C5 Minimum wage and robot adoption with a stable sample 166
C6 Minimum wage and robot adoption with lagged exposure in levels 167
C7 Minimum wage and robot adoption two-way clustered SEs (firm and year) 168
C8 Minimum wage and robot adoption without control variables 168
C9 Minimum wage and robot adoption with lagged control variables 169
C10 Robot adoption and firm-level outcomes (descriptive) 171
C11 Robot adoption, exposure (Patent-Based Occupational Exposure (PBOE)), and worker wages 178
C12 Robot adoption, exposure (PBOE), and worker displacement – fixed-term workers 179
C13 Robot adoption, exposure (PBOE), and worker displacement – all workers 180
C14 Robot adoption, exposure (PBOE), and worker reallocation – fixed-term workers 181
C15 Robot adoption, exposure (PBOE), and worker reallocation – all workers 182
C16 Robot adoption, exposure, and worker reallocation – any direction – fixed-term 183
C17 Robot adoption, exposure, and worker reallocation – any direction – all workers 184
C18 Robot adoption, exposure, and worker reallocation – any direction – open-ended 185

Acronyms

AI Artificial Intelligence *(pp. 7, 8, 79, 110)*

ATT Average Treatment on the Treated *(pp. 57, 65, 66, 136, 137, 155)*

CAE Portuguese Classification of Economic Activities *(pp. 124, 125)*

CEO Chief Executive Officer *(p. 15)*

CI International Trade Records *(pp. 1, 2, 44–46, 49, 51, 55, 67, 72, 74, 82, 84–86, 92–94, 96, 100, 101, 109, 150, 155–157, 159–161, 165–169, 171–177)*

CN Combined Nomenclature *(pp. 46, 49, 51, 54, 83, 109)*

CNC Computer Numerical Control *(pp. 46, 49, 51, 54)*

CPC Cooperative Patent Classification *(p. 81)*

CPP/2010 Classificação Portuguesa das Profissões 2010 *(p. 15)*

DiD Difference-in-Differences *(pp. xiv–xvi, 47, 50, 52, 56–58, 61, 62, 67, 68, 72, 73, 87, 132, 137, 138, 140, 145, 147, 159)*

EBIT Earnings Before Interest and Taxes *(p. 17)*

EU European Union *(p. 82)*

FMover First Mover *(pp. 23, 24)*

GPS Global Positioning System *(p. 14)*

GVA Gross Value Added *(pp. 26–32, 128–131)*

ICT Information and Communication Technologies *(pp. x, xii, 2, 6, 8, 9, 13, 14, 33, 35, 37, 39, 44–47, 52, 54–56, 58–70, 107–110, 136, 137, 155)*

INE Instituto Nacional de Estatística *(pp. 16, 84, 86, 92–94, 96, 100, 101, 165–169, 171–177)*

ISCO-08 International Standard Classification of Occupations (2008) *(pp. 83, 89)*

ISCO/2008 International Standard Classification of Occupations 2008 *(p. 15)*

IT Information Technology *(pp. 8, 9, 26)*

IUTICE Survey on Information and Communication Technologies in Enterprises *(pp. 1, 2, 4, 12–16, 18–20, 22, 24, 27, 28, 31, 32, 109, 110, 123, 124, 126, 128, 130, 131)*

LP Labour Productivity *(pp. 27, 28, 30–32, 129–131)*

MTSSS Ministry of Labour, Solidarity and Social Security *(p. 82)*

MW Minimum Wage *(pp. 75, 87)*

NACE Rev. 2 Statistical Classification of Economic Activities in the European Community (Revision 2) *(p. 82)*

NAdopter Non-Adopter *(pp. 23, 24)*

NUTS 2 Nomenclature of territorial units for statistics – level 2 *(pp. xii, xiv, 42–44, 52, 58, 159)*

O*NET Occupational Information Network *(p. 81)*

OLS Ordinary Least Squares *(p. 89)*

PBOE Patent-Based Occupational Exposure *(pp. xvi, 83, 89, 164, 178–182)*

PPML Poisson Pseudo-Maximum Likelihood *(pp. xv, 57, 136, 148)*

PRCI Potential Relative Cost Increase *(pp. 2, 79, 83, 84, 86, 87, 165–169)*

QP Personnel Records *(pp. 1, 2, 12–15, 18–20, 24, 27, 28, 31, 32, 44–46, 49, 51, 54, 61, 67, 72–74, 82, 83, 86, 87, 92–94, 96, 100, 101, 109, 123, 128, 130, 131, 133–135, 138–150, 155–157, 159–161, 164–169, 171–177)*

RTI Routine Task Index *(pp. 83, 89, 90, 95, 97, 164)*

SCIE Integrated Accounts System of Enterprises *(pp. xiv, 1, 2, 12–14, 18–20, 24, 27, 28, 31, 32, 44, 46, 47, 49, 51, 54–56, 68, 72, 73, 82, 85–87, 92–94, 96, 100, 101, 109, 123, 124, 126, 128, 130, 131, 133–135, 139, 164–169, 171–177)*

SITC Standard International Trade Classification *(p. 83)*

STAN A7 OECD Structural Analysis Database – A7 industry classification *(p. 82)*

STEM Science, Technology, Engineering and Mathematics *(p. 79)*

TFP Total Factor Productivity *(pp. 15, 20, 26–32, 80, 128–131)*

ths Thousands *(p. 86)*

TWFE Two-Way Fixed Effects *(pp. xv, 37, 52, 56, 57, 65, 66, 68, 87, 136–140, 142, 143, 145–147, 149, 150, 155–157)*

U.S. United States *(pp. 9, 78)*

1

Introduction

This thesis studies three margins of firm adjustment to technological and institutional change in the Portuguese economy: the timing of adoption of digital technologies, the recomposition of the workforce following a labour-supply shock, and the adoption of industrial robots following a labour-cost shock. The three empirical chapters share a common interpretive frame. Firms' responses are not uniform; they are mediated by internal capabilities (managerial and workforce education, organisational scale), by technology characteristics (modularity, complementarity with human capital), and by labour market conditions (the cost and availability of specific segments of the workforce). Across the three settings, the analysis combines administrative firm-level accounts, linked employer–employee records, and international trade data to trace how these three dimensions interact to shape firm behaviour and worker outcomes.[1]

The thesis addresses three questions. First, which firm-level characteristics predict the timing of adoption of digital technologies, and how does adoption timing map into firm performance – in particular, when do first-mover advantages persist rather than dissipate? Second, how do firms adjust when a binding institutional change restricts their access to a specific segment of the workforce (underage workers following the 2009 compulsory schooling reform) or raises the cost of a more pervasive segment (low-wage workers, following the statutory minimum wage increases of 2014–2021)? Third, within firms that adopt automation, how are the adjustment costs distributed across workers of different contractual status and different task exposure?

The empirical analysis draws on four administrative datasets provided by Statistics Portugal and the Ministry of Labour, Solidarity and Social Security: the SCIE, covering balance-sheet information for the near-universe of non-financial firms since 2006; Personnel Records (QP), the annual matched employer–employee register since 1986; the Survey on Information and Communication Technologies in Enterprises (IUTICE), the national component of the Eurostat ICT Usage Survey; and the International Trade Records (CI) trade microdata, used to identify firm-level imports of industrial robots through Combined Nomenclature following Acemoglu and Restrepo (2018a) and Acemoglu et al. (2023). Chapter 2 combines IUTICE, QP and SCIE across the 2014–2018 ICT waves and uses multi-sectoral data. Chapters 3 and 4 combine

[1] Chapter 2 is co-authored with Fernando Alexandre and Miguel Portela and is published as Alexandre et al. (2026). Chapter 3 is co-authored with Fernando Alexandre and Miguel Portela. Chapter 4 is co-authored with Hélder Costa.

QP, SCIE and CI and restrict the analysis to manufacturing firms. The linkage is implemented through a common anonymised firm identifier.

Chapter 2 studies the adoption of cloud computing and big data analytics as a function of firm-level capabilities, and the association between adoption timing and firm performance. Leveraging multiple waves of IUTICE survey, firms are classified as first movers, followers and non-adopters. The analysis then combines a multinomial logit with firm fixed-effects regressions and long-difference specifications, reweighted by a propensity score to mitigate selection bias. Firms with more educated managers and workers, and shorter managerial tenure, are more likely to adopt both technologies, especially big data. The performance implications are technology-specific. For cloud computing, early adoption yields gains in value added, turnover, employment and productivity, but followers catch up. For big data, early adoption yields productivity and value-added gains that do not dissipate over the sample period, and these gains are stronger when firms combine adoption with higher managerial and workforce education. The chapter interprets this contrast as evidence that a relatively standardised technology (cloud computing) generates transient first-mover rents, while a capability-intensive technology (big data) generates persistent advantages conditional on complementary human capital.

Chapter 3 studies whether Portugal's 2009 compulsory schooling reform, which raised the minimum school-leaving age from 15 to 18 (or to the completion of 12 years of education), induced differential firm adjustment in manufacturing. The reform is studied as a labour-supply restriction that removed access to a specific, low-cost input – underage workers – rather than as a general reduction in low-skilled labour supply. Although the share of minors was small in absolute terms (roughly 3–5% of the workforce in firms that employed at least one worker under age 18 before the reform), the results show that, in more exposed firms, this share declines by 2.2–3.2 percentage points, while the probability of employing any minor falls by 24–29 percentage points. The chapter further examines the reform impacts on firm-level robot imports, on an occupation-based proxy for automation-related workforce presence, and on complementary ICT employment. In a Difference-in-Differences design with firm fixed effects, the clearest differential post-reform response appears in ICT employment; the occupation-based proxy is also positive but imprecise, while annual robot import indicators remain close to zero. However, the ICT result is not reform-specific in a clean identification sense, as it is not robust across alternative placebo and specification checks. The contribution is therefore to document a differential post-reform pattern of firm adjustment with an explicit accounting of the design's vulnerabilities.

Chapter 4 studies whether statutory minimum wage increases induced firms to adopt industrial robots, and how robot adoption reshapes worker outcomes within adopting firms. The analysis covers Portuguese manufacturing firms over 2010–2021, a period during which the statutory minimum wage rose by roughly 30% in real terms. Firm exposure to the policy is measured by a predetermined version of the Potential Relative Cost Increase (PRCI) index of Alexandre et al. (2022), which captures the proportional wage-bill increase required to lift all sub-minimum workers up to the new floor; the predetermined formulation holds baseline wage structures fixed, so that variation reflects statutory change rather than endogenous wage adjustments. Using the heterogeneity-robust Difference-in-Differences estimator of Callaway and Sant'Anna

(2021) with continuous treatment intensity, the chapter documents that firms more exposed to minimum wage increases become more likely to import industrial robots, but that this response only materialises with a delay of three to five years, consistent with lumpy investment, planning and financing frictions. Conditional on adoption, the chapter turns to within-firm variation in workers' pre-adoption task exposure (measured by the routine-task-intensity index of Goos et al. (2014)) and in contract type. Among workers who remain in the same firm–occupation, those in more automatable tasks experience less favourable wage growth than less exposed co-workers within the same firm-year. Separations rise disproportionately among workers on fixed-term contracts with high task exposure, while observed internal reallocation for that group is limited. Open-ended contract workers, by contrast, adjust more frequently through internal reallocation to less exposed occupations.

Read together, the three chapters yield three cross-cutting findings. First, capabilities and technology complexity are jointly decisive for the returns to adoption: the returns to big data adoption depend on human capital, while the returns to cloud computing do not, because the two technologies differ in the intensity of complementary investments they require. Second, the response of firms to labour market pressures is not uniform, even when the pressure is binding. A narrow supply-side restriction (Chapter 3) operates primarily through workforce composition rather than capital adoption; a broader cost-side shock (Chapter 4) is instead associated with robot adoption at the extensive margin, but only after a lag of three to five years. The two results are consistent when viewed jointly: with a small affected input and fixed costs of capital adoption, workforce reshuffling dominates; with a pervasive cost shock, robot adoption becomes worthwhile, but frictions (given investment costs) delay it. Third, the incidence of automation is uneven within adopting firms: workers in more automatable tasks face wage penalties among stayers, and the probability of separation rises disproportionately for fixed-term workers. The incidence reflects not only occupational exposure, but also the institutional structure of the labour market, which concentrates the adjustment margin on the less protected contractual segment. The synthesis is therefore not that firms adjust through a single unified mechanism, but that the mediators of adjustment – capability, technology complexity, institutional segmentation, scale of the shock – can be identified and are consistent across settings.

The policy implications are chapter-specific but mutually reinforcing. Chapter 2 implies that the returns to public support for digital adoption depend on complementary investments in managerial and workforce education; absent such complementarities, subsidised adoption of complex technologies such as big data will produce weaker productivity gains. Chapter 3 implies that compulsory schooling reforms have identifiable firm-side consequences, but that the evidence for reform-induced automation – as opposed to reform-induced workforce recomposition – is not robust. Chapter 4 implies that evaluations of minimum wage policy that consider only contemporaneous employment effects miss a medium-run technological-substitution margin, and that, in segmented labour markets, the incidence of adjustment concentrates on fixed-term workers. Taken together, the three chapters support coordinating education, wage and innovation policies, and in particular policies that support skill formation and transitions for workers in the most exposed segments of the workforce.

The findings are subject to three scope restrictions. First, Chapters 3 and 4 focus on the Portuguese manufacturing sector; Chapter 2 is multi-sectoral but restricted to the IUTICE sample. Second, robot adoption is measured through international trade records, which capture direct imports but not domestically produced or indirectly acquired robots; the resulting indicator is a lower bound. Third, the IUTICE coverage of cloud computing and big data is intermittent, which limits panel-based inference in Chapter 2. These restrictions are explicit and do not undermine the internal validity of the design within each chapter.

The remainder of the thesis is organised as follows. Chapter 2 studies the timing of adoption of cloud computing and big data and the associated performance outcomes. Chapter 3 studies the firm-level consequences of the 2009 compulsory schooling reform. Chapter 4 studies the effect of statutory minimum wage increases on firm-level robot adoption and the within-firm distribution of worker outcomes following adoption. Chapter 5 synthesises the findings and discusses their policy implications.

2

Timing of cloud computing and big data adoption: Cohort characteristics and firm performance

2.1 Introduction

This paper examines which firm characteristics influence the adoption of two digital technologies – cloud computing and big data – and how the timing of adoption affects subsequent firm performance.[1] Theoretical frameworks on technology diffusion and competitive advantage suggest that adopters may secure productivity (Bartelsman, 2013) and efficiency (Bartel et al., 2007) gains by leveraging innovative technologies. However, empirical evidence on the relationship between adoption timing and firm performance remains limited. To address this gap, we focus on two research questions. First, what firm-level characteristics differentiate first movers, followers, and non-adopters of cloud computing and big data? Second, do early adopters of these technologies exhibit superior performance – such as productivity – compared to firms that adopt later or the ones that do not adopt?

Digitalization has significantly reshaped the business landscape, triggering profound changes in processes and markets. While evidence suggests that digital adoption enhances firm-level productivity (Cette et al., 2022; Hall et al., 2013; Koch et al., 2021), considerable dispersion exists across firms (Bresnahan et al., 2002; L. Wu et al., 2020), with uncertainties prevailing at the aggregate level (Acemoglu et al., 2014; Fernald et al., 2024). Understanding the factors contributing to maximizing the benefits of these technologies is paramount. The interplay of various factors influences the outcomes of digital adoption and the subsequent impacts on firm performance (Calvino & Fontanelli, 2023). Complementary innovations, investments in intangible capital, organizational changes, and developing essential skills are among the crucial factors (Berlingieri et al., 2020; Brynjolfsson et al., 2019; Calvino et al., 2022). Management characteristics also play a pivotal role in facilitating technology adoption and enhancing productivity through innovative processes and dynamic organizational cultures (Bloom, Eifert, et al., 2012; Calvino & Fontanelli, 2023; Giorcelli, 2019; Koch et al., 2021). Furthermore, the timing of technology adoption emerges as a critical factor for competitiveness (Acemoglu et al., 2023) and future innovation (Engelstätter, 2012;

[1] This chapter is co-authored with Fernando Alexandre and Miguel Portela and is published as Alexandre et al. (2026).

Niebel et al., 2019). Nevertheless, embracing digital technology as an early adopter entails inherent risks stemming from uncertainty and unpredictability. Innovators invest significant resources without historical data, rendering them vulnerable to unforeseen challenges such as technical glitches and compatibility issues. Therefore, comprehending the performance impacts and determinants of technological adoption across various cohorts is essential for identifying potential competitive advantages for early adopters.

To conduct our analysis we use a comprehensive longitudinal sample of Portuguese firms and focus exclusively on cloud computing and big data analytics; our results should not be extrapolated mechanically to other technologies or countries. The study combines annual waves of the national ICT survey with matched balance-sheet records and linked employer–employee data for the population of firms, yielding rich information on technology uptake, management attributes, workforce composition and economic outcomes. Throughout, the term "digital technologies" refers solely to cloud computing and big data analytics.

Our empirical strategy unfolds in four steps. Firstly, we allocate firms to the cohorts first mover, follower, and non-adopter, describe adoption shares across industries and summarise cohort-specific firm characteristics. Secondly, a multinomial logit estimates how management and workforce variables predict cohort membership. Thirdly, to estimate the performance effect of being an early adopter versus non-adopter, we run firm-fixed-effects regressions combined with propensity-score weights to mitigate selection bias. Finally, a long-difference specification estimation contrasts first movers with followers, isolating the incremental advantage (or disadvantage) of pioneering adoption.

Our findings are fourfold: (i) Firms led by more highly educated managers are more likely to be first movers in the adoption of cloud computing and big data technologies, whereas older firms and those with longer-tenured managers are less likely to take the lead in adopting these innovations; (ii) Early adopters of cloud computing and big data experience significant gains in productivity and performance relative to non-adopters, with the benefits being somewhat more pronounced among cloud computing adopters; (iii) The positive effect of big data adoption on firm productivity is significantly amplified by higher levels of education among both managers and employees; (iv) In contrast to cloud computing – where late adopters can more readily close the performance gap – first movers in big data maintain a persistent advantage in productivity and gross value added over subsequent adopters.

The remainder of this paper is organized as follows. Section 2.2 reviews the literature on technology adoption, firm performance, aggregate productivity, management quality and adoption timing. Section 2.3 presents the conceptual framework. Section 2.4 describes the data, variables and descriptive statistics. The methodology is outlined in Section 2.5. Section 2.6 presents and discusses the results. Finally, Section 2.7 concludes by summarizing key findings and suggesting implications for policymakers and firm managers.

2.2 Literature review

The global diffusion of the Internet since the 1990s triggered a far-reaching digital transformation, enabling technologies such as cloud computing and big data and opening more dynamic yet volatile markets. Successful uptake of these general-purpose technologies depends on broader social readiness (Brynjolfsson & McAfee, 2014) – including the alignment of workers' skills with complementary tools (Brynjolfsson & McAfee, 2012), organisational redesign and decentralisation (Bresnahan et al., 2002), and competitive market conditions that foster diffusion (Andrews et al., 2016).

To motivate our study, the review proceeds in four subsections: (i) characteristics of technology adopters, (ii) impacts on firm performance, (iii) the role of management, and (iv) the significance of adoption timing.

2.2.1 Technology adoption

Technology adoption is uneven across firms. Larger firms, benefiting from economies of scale, show a higher propensity to introduce new digital tools (Conti et al., 2023; Haller & Siedschlag, 2011; Koch et al., 2021). Portuguese evidence confirms that big data analytics is primarily adopted by large firms to streamline processes and cut costs, whereas smaller firms deploy it mainly to support product innovation (Conti et al., 2023). This size-dependent pattern reflects divergent ways firms leverage technology: large firms prioritise operational efficiency, while small firms seek revenue growth through new offerings.

Export-oriented firms tend to adopt digital technologies earlier, responding to intense international competition and the need for productivity gains (Calvino & Fontanelli, 2023; Haller & Siedschlag, 2011; Koch et al., 2021). Spillover effects act as a powerful catalyst for innovation: they pressure non-adopters to match the technologies of leading firms in order to stay competitive and avoid losing market power, thereby helping diffuse the efficiency gains realised by early adopters (Acemoglu et al., 2023; Engelstätter, 2012; Haller & Siedschlag, 2011). Within the group of adopters, technological complementarities reinforce this dynamic, as firms frequently implement cloud, big data, Artificial Intelligence (AI), and related tools together to maximise joint benefits (Cho et al., 2023).

The relationship between worker skills and technology adoption runs in both directions. Firms that adopt digital technologies typically employ more highly skilled workers, and the ensuing productivity gains are most pronounced for these employees in micro and small enterprises (Calvino et al., 2022). Consistently, Barbosa and Faria (2022) stress that systematic upskilling is essential to realise the full benefits of new technologies.

Unequal access to advanced technologies widens productivity and market-power gaps across firms. Frontier firms concentrate cutting-edge tools and skilled labour, reinforcing their dominance, while laggards fall further behind. Andrews et al. (2016) show that much of the recent slowdown in aggregate productivity stems from this growing gap: leaders continue to post strong gains, whereas other firms face diffusion barriers that hinder adoption. Even young firms with qualified workers struggle to catch up in highly

digital sectors, highlighting frictions in technology and knowledge transfer (Berlingieri et al., 2020). Policy measures that ease diffusion and bolster workforce skills across the firm distribution are therefore critical to spread digital gains and raise overall productivity (Andrews et al., 2015).

2.2.2 Impacts on firm performance and aggregate productivity

Evidence on how technology adoption affects firm performance has evolved substantially. Early studies, such as Stoneman and Kwon (1996), found that new technologies could raise firm profits by about 11%, although the size of the gain varied widely across firms and industries. Subsequent work confirms the positive association but shows that its magnitude depends on several complementary factors, including management quality (Bloom & Van Reenen, 2007), firm size (Bartelsman, 2013; T.-H. Wu et al., 2022), the composition of human capital (Bartel et al., 2007), firms' intangible–asset stock (Bartelsman, 2013; T.-H. Wu et al., 2022), and opportunities for spillover learning (Segarra-Blasco et al., 2022).

Recent evidence for French, Italian, and Portuguese firms shows a clear positive link between ICT adoption and productivity (Amador & Silva, 2023; Cette et al., 2022; Hall et al., 2013). For cloud computing the picture is mixed: some studies report sizeable productivity gains (Cette et al., 2022; Duso & Schiersch, 2022), whereas others find no statistically significant effects (Gal et al., 2019). Impacts also vary by firm type: young adopters often enjoy strong employment and revenue growth, while mature incumbents see more modest gains (DeStefano et al., 2023).

Big data adoption raises productivity by roughly 3–7 % on average, though benefits are concentrated in IT-intensive or highly competitive industries (Müller et al., 2018). Brynjolfsson, Jin, and McElheran (2021) show that predictive analytics boosts performance only when accompanied by complementary assets such as Information Technology (IT) capital, skilled labour, and streamlined workflows – explaining why some adopters see limited pay-offs. Survey evidence for Germany likewise finds a positive association between AI adoption and productivity (Czarnitzki et al., 2023). Firms that integrate data analytics into process improvement and innovation capture larger gains than those using the technology for standalone development projects (L. Wu et al., 2020). Finally, higher intangible-asset stocks – especially digital-skill intensity – amplify productivity effects, with pronounced advantages in the service sector and among younger firms (Borowiecki et al., 2021).

Despite firm-level evidence of sizeable gains from early technology adoption, aggregate productivity performance remains disappointing. ICT investment was a major driver of the 1990s productivity surge (Stiroh, 2002), yet the hoped-for sustained acceleration has not materialised, even as digital tools have spread widely in the twenty-first century. Acemoglu et al. (2014) argue that some early gains reflected labour displacement rather than genuine value added, especially in computer-intensive manufacturing sectors. Likewise, Fernald et al. (2024) link the mid-2000s slowdown in advanced economies to weaker total-factor-productivity growth, challenging the assertion that continued ICT adoption would sustain earlier productivity gains. They suggest that ICT acted as a general-purpose technology that drove growth only

temporarily before its transformative effects diminished. A further complication is measurement: intangible investments, quality improvements, and consumer surplus generated by digital technologies are only partially captured by standard metrics (Brynjolfsson, Rock, & Syverson, 2021; Cette & Spiezia, 2022). The resulting "dark matter" of unmeasured gains underscores the limits of conventional productivity statistics for assessing the full economic impact of technology adoption (Brynjolfsson, Rock, & Syverson, 2021).

Finally, it has been argued that the New Digital Economy remains in its "installation phase", during which firms invest heavily in digital infrastructure and experiment with new technologies but realise only limited productivity pay-offs. Van Ark (2016) notes that this stage is marked by sizeable ICT spending yet little immediate, broad-based productivity growth. Meaningful gains are expected to emerge only in the subsequent "deployment phase", once the technologies diffuse widely, become fully embedded in business processes, and propagate across industries – conditions under which measurable productivity improvements should materialise.

These studies highlight the subtle, multifaceted link between technology adoption and performance, underscoring the need to account for the many factors that mediate that relationship.

2.2.3 Role of management

Well-managed firms tend to be more efficient, grow faster, and show greater resilience (Alexandre et al., 2021; Bloom, Eifert, et al., 2012). Furthermore, management quality is a decisive factor in both the adoption and the effective use of new technologies, with far-reaching consequences for firm performance and competitive advantage (Bloom, Sadun, & Reenen, 2012; Bloom & Van Reenen, 2007; Bloom et al., 2019; Daveri & Parisi, 2015). This advantage stems from what skilled managers actually do: they (i) recognise technological opportunities, (ii) mobilise the requisite intangible- and human-capital resources, and (iii) steer the organisation through the complexities of technology integration (Bartelsman, 2013; Bloom et al., 2019; Bresnahan et al., 2002).

Empirical evidence confirms these channels. For example, Bloom, Sadun, and Reenen (2012) show that United States (U.S.) multinationals operating in the United Kingdom are markedly more productive than comparable non-U.S. firms. They trace this gap to superior management practices that create more effective organisational structures and allow fuller exploitation of new technologies.

Moreover, stronger management capabilities amplify – and help sustain – the benefits of technology adoption over time (Giorcelli, 2019; T.-H. Wu et al., 2022). For instance, Bartel et al. (2007) show that introducing IT-enhanced machinery prompts workplace reorganisation and streamlines managerial processes, boosting efficiency throughout the production cycle.

At the macro level, economies in which firms pair widespread IT adoption with advanced management practices record faster productivity and income growth (Schivardi & Schmitz, 2019). Complementing this aggregate evidence, managerial demographics also matter: Daveri and Parisi (2015) show that older managers can dampen firm performance, particularly in innovation-intensive settings.

In short, capable management unlocks the full value of digital technologies, translating adoption into sustained innovation, efficiency, and competitive advantage.

2.2.4 Significance of timing

The literature explicitly addressing the timing of technology adoption – comparing firms that adopt earlier versus later – is relatively scarce. Still, some studies show that adopting one technology can facilitate subsequent innovation. Niebel et al. (2019), for instance, show that embracing big data raises both the likelihood and intensity of later innovation. In a similar vein, Engelstätter (2012) show that implementing enterprise software systems increases the likelihood of future innovation. Cho et al. (2023) further emphasise complementarities across technologies that generate, process, and manage large volumes of data – such as big data, the Internet of Things, cloud computing, and artificial intelligence – which in turn support the adoption of other advanced tools like 3D printing and robotics.

While some studies link early technology adoption to increased innovation potential, evidence on how adoption timing affects firm-level performance, namely productivity, remains limited. The available research suggests that firms adopting later may struggle to fully capitalise on future innovations, in part because technological capabilities often build cumulatively (Cho et al., 2023; Engelstätter, 2012; Niebel et al., 2019). As discussed in Section 2.2.1, early adopters tend to be systematically different – larger, better-resourced, and with more educated workforces – highlighting the existence of structural barriers to adoption. These disparities suggest that timing not only shapes firms' capacity to innovate, but also affects their ability to translate digital investments into productivity gains.

Timing also shapes competitive dynamics. Acemoglu et al. (2023) report that firms slow to adopt digital technologies experience weaker catch-up rates in industries that rely heavily on digital skills. Thus late movers risk persistent productivity gaps and loss of market share.

Taken together, early adoption of digital technologies may confer a cumulative advantage: by innovating sooner, firms lock in gains that compound over time, widening the gap with rivals. Because such early uptake remains concentrated among a relatively small set of well-positioned firms, non-adopters face mounting hurdles in narrowing that distance, so performance differentials are more likely to widen than to close. Yet this prospective edge is tempered by first-mover risk: pioneers commit resources before reliable benchmarks exist and thus expose themselves to technical failures and compatibility problems.

2.3 Conceptual framework and hypotheses

To guide the empirical analysis and clarify the mechanisms linking firm characteristics, technology adoption timing, and performance outcomes, we propose a conceptual framework grounded in established economic literature. This framework integrates three dimensions: (i) internal firm characteristics influencing adoption decisions and timing for cloud computing and big data; (ii) adoption timing categories

(first movers, followers, non-adopters); and (iii) performance outcomes moderated by human capital and technology characteristics.

Four interconnected mechanisms drive technology adoption and its performance effects. First, internal capabilities – particularly managerial and worker education – determine firms' capacity to evaluate technologies, mobilize resources, and implement organizational change (Bloom & Van Reenen, 2007; Calvino et al., 2022; Daveri & Parisi, 2015). Higher educational endowments enhance absorptive capacity, reduce implementation costs, and enable recognition of technological opportunities.

Second, organizational characteristics create path dependencies shaping adoption decisions. Firm age and managerial tenure generate accumulated routines that provide assessment knowledge but simultaneously create resistance to change (Daveri & Parisi, 2015). This tension between experience-based capabilities and organizational inertia influences both the likelihood and timing of adoption.

Third, technology characteristics determine requisite capabilities and implementation complexity. Cloud computing represents a standardized, modular technology with vendor support and minimal customization requirements. Big data analytics demands substantial complementary investments in analytical capabilities, software infrastructure, and organizational processes (Brynjolfsson, Jin, & McElheran, 2021; Müller et al., 2018). These complexity differences fundamentally shape adoption patterns and performance outcomes.

Fourth, these mechanisms interact to determine adoption outcomes. The effect of human capital on adoption is moderated by organizational characteristics – younger firms leverage educated managers more effectively due to lower structural rigidity. Performance benefits from early adoption materialize only when sufficient human capital exists to exploit the technology. This interaction intensifies with technology complexity: for big data, the complementarity between timing and capabilities creates a separating equilibrium where only high-capability firms benefit from early adoption, while for cloud computing, standardization reduces these complementarities (Engelstätter, 2012; Niebel et al., 2019).

Based on these mechanisms, we formulate four testable hypotheses:

- Hypothesis 1: Firms with higher managerial and worker education levels exhibit greater probability of adopting cloud computing and big data.

- Hypothesis 2: Firm age and managerial tenure negatively associate with early adoption probability, reflecting the dominance of inertia effects over experience benefits.

- Hypothesis 3: The relationship between human capital and both adoption probability and performance gains is stronger for big data than cloud computing, reflecting differential complexity and complementarity requirements.

- Hypothesis 4: First-mover productivity advantages are insignificant for cloud computing but significant and persistent for big data, due to differential learning curves and capability-technology complementarities.

These hypotheses extend existing theory by proposing that technology characteristics moderate the relationships between capabilities, timing, and performance – challenging the implicit assumption of homogeneous digital technologies in prior literature. While Hypotheses 1 and 2 test established predictions, Hypotheses 3 and 4 offer novel insights: the differential importance of human capital across technologies and the technology-contingent nature of first-mover advantages.

This framework predicts that standardized technologies (cloud computing) exhibit rapid diffusion eliminating first-mover advantages, while complex technologies (big data) generate persistent competitive advantages through cumulative learning. Understanding these heterogeneous effects contributes to explaining aggregate productivity implications of digital transformation.

The framework directly informs our empirical strategy. Section 2.5 employs multinomial logit models to test how firm characteristics predict adoption timing (Hypotheses 1-2). Fixed-effects and long-difference specifications with propensity-score reweighting estimate differential performance effects and human capital complementarities across technologies (Hypotheses 3-4), enabling identification of whether theorized mechanisms operate as predicted.

2.4 Data and variables description

This section provides a description of the database and a discussion of the descriptive statistics.

2.4.1 Data

The data used in this study comes from three primary sources. The SCIE (*Sistema de Contas Integradas das Empresas*), the QP (*Quadros de Pessoal*), and the IUTICE, all provided by Statistics Portugal. These three datasets are linked through a common anonymized identifier.[2]

The SCIE dataset, INE, 2021b, includes comprehensive information on firms' balance sheets and financial statements for all non-financial firms starting from 2006 (about 350,000 firms per year). It results from integrating statistical information on companies, primarily based on administrative data, focusing on simplified business information. This dataset allows for the computation of firm-level data on economic and financial indicators.

The QP dataset provides information on the Portuguese labour force and has been collected yearly since 1986 by the Ministry of Labour, Solidarity, and Social Security (INE, 2021c). It covers all private firms with at least one employee, its establishments, and employees in Portugal, amounting to about 3 million workers and 350 thousand firms annually. This dataset offers detailed worker-level information on various employee characteristics, such as tenure, educational level, monthly earnings, hours worked, occupation, and identification of the employing firm.

The central database used in this study is IUTICE, an annual stratified probability sample survey collected by Statistics Portugal (INE, 2021a). Implemented since 2004, the IUTICE focuses on firms adopting

[2]See Figure 12 in the Appendix for a flowchart illustrating the database construction process.

information and communication technologies and e-commerce. It is part of the Eurostat Community Survey on ICT Usage and e-commerce in Enterprises. This survey aims to enhance our comprehension of the adoption of information and communication technologies in the business sector and examine the significance and prioritization of these technologies in relation to overall economic competitiveness.

The IUTICE survey gathers data from a representative sample of companies in Portugal operating in manufacturing, energy, construction, trade and repair, accommodation and catering, transport and communications, and other services (excluding education, health activities, and financial activities). Large firms with over 250 workers or total revenues exceeding 25 million euros are surveyed comprehensively. In contrast, smaller firms are sampled using a stratified random approach based on the number of workers, turnover, and economic activity. Sampling weights are provided in the survey to account for the over-representation of large-sized firms. Participation in the IUTICE survey is mandatory for selected firms. IUTICE is an unbalanced dataset with approximately seven thousand firms each year.

IUTICE has three distinct weights: one based on the number of firms, another on the number of workers, and a third on turnover. According to the survey's methodology, the weight based on the number of firms should be assigned to most of the variables considered. This weight is determined by the inverse probability of a company being selected for inclusion in the sample. The second weight is designed for variables related to the number of workers, while the third pertains to the variables associated with turnover. As discussed in Heeringa et al. (2017), the sample selection weight ensures that our sample accurately represents the target population. Unbiased estimation of the population distribution requires the weighting of sample data when population elements have varying probabilities of being included in the sample, as is the case here (Heeringa et al., 2017).

Therefore, in this study, all inferences and estimations are derived using sample survey weights. The weight for the number of firms is used for descriptive statistics and the analysis of adoption probabilities. The turnover weight is employed to assess the impacts of technology on productivity and turnover. Additionally, the weight of the number of workers is applied to examine the effects of technology on employment.

In this study, we focus on two digital technologies: cloud computing, surveyed in 2014, 2016, 2017, and 2018, and big data, surveyed in 2016 and 2018.[3] The initial IUTICE dataset comprises 39,411 observations from 26,018 firms, covering the period from 2014 to 2019. To align with the objectives of our study, we excluded all observations for which firms reported not using the internet in the corresponding year, resulting in the removal of 3,821 observations, corresponding to 3,458 unique firms. This left us with 35,590 observations from 22,805 firms. We then merged the IUTICE data with the QP and SCIE datasets, which led to a further reduction of 4,202 observations: 497 were missing in both QP and SCIE, and 3,705 were present only in QP. Our final sample comprises 19,828 firms. For these firms, we collected data from QP and SCIE covering the years 2006 to 2019. The resulting firm-level panel comprises 238,717

[3]This intermittent survey coverage, combined with the sampling design of IUTICE for smaller firms, limits the longitudinal dimension of the data. Because of these constraints, we cannot fully leverage typical panel-data techniques. Instead, we adopt the identification strategies described in Section 4 that rely principally on cross-sectional estimators.

observations.[4]

2.4.2 Variable description

Cloud computing refers to internet-based ICT services that provide access to software, storage capacity, and computing power, among other features. It operates through a global network of servers located worldwide. These servers are designed to manage data and enable the execution of applications and content accessible from anywhere, anytime, via the Internet. Cloud computing offers advanced security and file recovery capabilities. Firms can customize their usage, acquiring the necessary space and functionality without additional hardware investments. This empowers firms of all sizes to securely store and manage files, data, documents, and applications, resulting in faster access than traditional reliance on personal computers and manual file transfers. Cloud computing enhances management efficiency, protects firm information, and enables seamless adaptability to expansion, contraction, or new international strategies without the need for complex hardware investment projects. In this study, cloud computing includes office software, file and database storage, accounting or finance applications, customer relationship management, and running their software (excluding firms that solely utilize cloud computing for email services).

Big data refers to information gathered from firm-owned sensors or smart devices, geolocation data from portable devices, data generated from digital media (e.g., social media), and other unspecified sources. Big data is characterized by three main attributes: volume (large quantities of data generated over time), variety (diverse formats of complex data, both structured and unstructured, such as text, videos, images, voice, documents, sensor data, etc.), and velocity (the rapid pace at which data is generated and changes over time). Firms employ big data technologies to collect, process, and analyze data to identify patterns, trends, and insights. This process involves specialized software tools, algorithms, and analytical skills to manage and interpret the data. Big data can enhance operational efficiency, customer targeting, decision-making processes, and other aspects. Examples of big data applications include targeted retail offers, Global Positioning System (GPS) navigation assistance, and business intelligence. It is worth noting that big data often requires significant data storage and processing capabilities, frequently leveraging cloud computing infrastructure.

We define first movers as those firms that adopted each technology in its inaugural year of inclusion in the IUTICE survey: cloud computing was included in 2014, and big data in 2016. However, we should note

[4] As pointed out by an anonymous referee, we experienced a relevant loss of observations when merging with the other datasets, QP and SCIE, raising concerns about non-random sample attrition. To address this issue, we re-evaluated our results by implementing a correction to the weights. This correction involved stratifying the population based on three key characteristics: primary sector of activity, workforce size category, and turnover category. Each stratum comprises firms sharing the same sector, workforce size, and turnover category. We then recalculated the weights for each stratum – accounting for the number of firms, workforce size, and gross value added – to address the observation loss. These recalculated weights were subsequently redistributed evenly among the remaining firms within the same stratum to maintain alignment with the original dataset. The main results of our paper discussed in Section 5 remain consistent; we observed enhanced significance in certain estimators and slight variations in the magnitude of the estimated impacts compared to our initial findings. Results are available upon request.

that these technologies could have been already in use in previous years, which limits the precision of our definition of first movers. Nevertheless, considering the relatively low adoption rates among first movers (6.76% for cloud computing and 9.67% for big data) and Portugal's non-pioneering position in technology adoption, we believe our definition is adequate to identify first movers. Follower firms are identified as those that adopt the technology in the subsequent survey waves. Non-adopter firms are those that have never reported the usage of the technology in all waves of the survey.

Another dimension relevant to our analysis is the role of the managerial team's competencies in decisions to adopt cloud computing and big data technologies. Executive decisions are not made solely by the general manager or Chief Executive Officer (CEO) but involve managers reporting to the administration board as high-level directors or key senior managers (Bloom & Reenen, 2011; Carpenter, 2002). To capture this, we constructed the variable "Manager," which includes all top managers–namely, CEOs, executive directors, and the highest-level department directors in firms. This variable is based on data from the Personnel Records data QP regarding workers' occupations, categorized according to the Portuguese Classification of Occupations 2010 (Classificação Portuguesa das Profissões 2010 (CPP/2010)) and aligned with the International Standard Classification of Occupations 2008 (ISCO/2008). The management team includes managing directors and chief executives, administrative and commercial managers, production and specialized services managers, as well as hospitality, retail, and other services managers.

Concerning the impact of digital technologies on performance, we consider the following variables: gross value added (GVA), turnover, employment, labour productivity (LP), measured by the ratio of value added over the number of workers, and total factor productivity (TFP). TFP is computed following Ackerberg et al. (2015).[5]

Table 31 in the Appendix comprehensively describes all variables used in this study.[6]

2.4.3 Descriptive statistics

This section discusses descriptive statistics for the variables underlying our empirical analysis. Table 2 presents the adoption rates for each technology, categorized by firm size. In 2014, the initial reporting year for cloud computing, 6.8% of firms adopted that technology. Adoption rates increase with firm size. The digital adoption rate was 5.7% for micro-sized firms, 7.7% for small-sized firms, 10.5% for medium-sized firms, and 15.2% for large-sized firms. In 2016, big data technology was adopted by 9.7% of the firms in the IUTICE survey. As in cloud computing technology, the big data adoption rates increase with firm size: 8.2% for micro-sized firms, 11.2% for small-sized firms, 15.4% for medium-sized firms, and 20.6% for large-sized firms. This pattern is observed in all waves of the survey. This finding aligns with existing

[5] Total Factor Productivity (TFP) was calculated using the Stata software with the "acfest" command.

[6] For classification, we define micro firms as those with fewer than 10 workers and turnover and total assets equal to or less than 2 Million Euros. Small firms have less than 50 employees (with a minimum of 10) and turnover and total assets equal to or less than 10 Million Euros. Medium-sized firms employ between 50 and 250 workers, with turnover ranging from 10 Million to 50 Million Euros and total assets between 10 Million and 43 Million Euros. Large firms have more than 250 workers, turnover exceeding 50 Million Euros, and total assets surpassing 43 Million Euros.

literature, highlighting the positive correlation between technology adoption and firm size (Berlingieri et al., 2020; Calvino & Fontanelli, 2023; Koch et al., 2021).

Table 2 also shows increasing adoption rates of cloud computing over time, overall and by firm size. For the sample including all firms, the incidence of cloud computing increased from 6.7% in 2014 to 14.5% in 2018. For the population of large firms, there was an increase from 15.2% in 2014 to 47.7% in 2018. However, the upward trend over time, except for large-size firms, is not observed for big data technology. Nevertheless, while the adoption rate for big data decreases for micro-, small-, and medium-sized firms, it shows a significant increase (by 9.3 percentage points) for large-sized firms from 2016 to 2018.[7]

Table 2: Percentage of firms that have adopted technology, broken down by size and year (%)

	2014	2016	2017	2018
	Panel A: all firms			
Cloud computing	6.76	10.42	12.14	14.53
Big data		9.67		8.53
	Panel B: micro firms			
Cloud computing	5.74	8.13	9.42	10.99
Big data		8.22		6.28
	Panel C: small firms			
Cloud computing	7.65	12.92	13.52	17.34
Big data		11.23		10.97
	Panel D: medium firms			
Cloud computing	10.50	19.32	25.19	27.76
Big data		15.42		14.08
	Panel E: large firms			
Cloud computing	15.24	27.57	41.96	47.74
Big data		20.60		29.90
Number of firms	4,584	5,238	5,549	5,422

Notes: We report the percentage of adopters of each technology among the groups of firms divided by size in panels B to E. Cloud computing technology is surveyed in 2014, 2016, 2017, and 2018; and Big data in 2016 and 2018. This report uses sampling weights provided by the survey. *Source:* Authors' calculations based on IUTICE INE, 2021a.

When analyzing the adoption rate of cloud computing by industry, our data show that the "information

[7] An anonymous referee noted that overall cloud adoption appears relatively low, whereas big data adoption seems comparatively high. We emphasize that our figures are not always directly comparable to those in other reports or studies. For example, some studies (e.g., Conti et al., 2023) do not apply survey weights, while others restrict their samples to a subset of firms. In particular, Statistics Portugal (INE) excludes micro-sized firms from its official reports to align with international standards, as many countries lack data on very small firms. When we apply a similar exclusion, our estimates of big data adoption align closely with the upward trend reported by Instituto Nacional de Estatística (INE) for larger firms. However, INE has confirmed that micro firms are accurately represented in IUTICE. Given that they constitute the majority of Portuguese firms, we chose to include them in our analysis.

and communication services sector" has the highest adoption rate.[8]. For example, in 2018, 46.0% of firms in that sector adopted cloud computing. That adoption rate was closely followed by the "other services: repair of computers and communication equipment sector", with an adoption rate of 34.2%, and the "water, sanitation, waste management and depollution", which recorded a rate of 30.5%. Moreover, when considering only the firms that adopted cloud computing, the largest share (26.6%) came from the "wholesale and retail sector", primarily from the "wholesale sector". "Manufacturing industries" ranked as the third-highest sector in terms of cloud computing adopters, accounting for 15.3%, preceded by "consulting, scientific, technical, and similar services", with an adoption rate of 18.6%.

Concerning big data technology, the sectors with the highest adoption rates in 2018 were "water, sanitation, waste management, and depollution" (20.9%), "electricity, gas, steam, and cold air" (20.8%), and "information and communication services" (20.6%). When examining the distribution of big data adopters by sector, the largest percentage of adopters, comprising 29.3%, was again observed in the "wholesale and retail sector" (specifically in the retail sub-sector). This was followed by "manufacturing industries" (17.1%) and "consulting, scientific, technical, and similar services" (12.9%).

Tables 3 and 4 provide descriptive statistics for cloud computing and big data, respectively. These tables classify firms into non-adopters, first movers, and followers.

Concerning cloud computing, Table 3, we observe that non-adopters have the lowest mean investment in fixed assets (77 thousand euros) compared to first movers (323 thousand euros) and followers (363 thousand euros). The education level of managers and the workforce, measured in years of education, is consistent across the three groups, with mean values around 11 to 13. However, first movers and followers have a higher proportion of highly educated managers and workers than non-adopters, suggesting that firms adopting cloud computing tend to employ a more educated workforce. Managers' tenure is quite similar across the three categories, with only slight variations in mean values. Non-adopters have the lowest percentage of foreign-owned firms (3%), while followers have the highest (7%). First movers and followers have a significantly higher mean number of workers (48 and 54, respectively) than non-adopters (15). Both first movers and followers exhibit a substantially higher mean turnover (7,771 thousand euros and 9,692 thousand euros) and gross value added (1,900 thousand euros and 2,141 thousand euros) compared to non-adopters (1,915 in turnover and 440 in gross value added), indicating that adopting firms are larger. First movers and followers also have a higher percentage of computer users than non-adopters, with mean values of 68% and 73%, respectively.[9]

Regarding big data technology, Table 4, it is worth noting that non-adopters have a lower mean investment in fixed assets (81 thousand euros) compared to first movers (257 thousand euros) and followers (515 thousand euros). However, it is interesting to observe that the median values for all three categories

[8]For a detailed analysis, see Appendix A.1

[9]Note: The extreme negative mean profitability for Followers in the Cloud Computing group is driven by a single outlier firm with an exceptionally low Earnings Before Interest and Taxes (EBIT). Since this variable is not used in the main regressions, we report it unadjusted, with the median serving as a more representative measure.

Table 3: Descriptive statistics – Cloud Computing

	Non adopters		First Movers		Followers	
	Mean	P50	Mean	P50	Mean	P50
Investments in fixed assets (in ths)	77	2	323	2	363	9
Firm age	19	16	17	13	19	16
Average wage (in ths)	16	13	18	15	21	17
Capital intensity (in ths)	46	7	62	6	43	7
Education of managers	11	11	13	13	13	13
Education of workforce	10	9	11	11	11	11
Exporter (%)	7	0	7	0	9	0
Foreign-owned (%)	3	0	6	0	7	0
Leverage (%)	84	62	72	66	86	63
Previous management experience	0	0	0	0	1	0
Number of workers	15	6	48	7	54	9
Productivity (in ths)	29	18	35	20	37	23
Profitability (%)	2	4	3	4	-42	5
Share high educ. managers (%)	25	0	46	50	49	50
Share high educ workers (%)	10	0	20	8	22	11
Tenure managers	13	12	11	10	12	11
Turnover (in ths)	1915	338	7711	487	9692	591
Gross value added (in ths)	440	100	1900	160	2141	206
Share of computer users (%)	60	56	68	75	73	88

Notes: We report all years of the sample. Sampling weights provided by the survey are used in this report. P50 stands for percentile 50. "ths" stands for thousands. *Source:* Authors' calculations based on IUTICE, QP and SCIE (INE, 2021a, INE, 2021c, and INE, 2021b).

are relatively low, suggesting a significant variation in asset investments within each group. The education level of managers and the workforce shows minimal variation across the three groups, with mean values around 10 to 12 years of schooling. First movers and followers have a higher proportion of highly educated managers and workers compared to non-adopters. The share of exports on total turnover is slightly higher for first movers (10%) and followers (10%) compared to non-adopters (7%). Non-adopters have the lowest percentage of foreign-owned firms (3%), while followers have the highest (7%). Followers also have the highest mean number of workers (63) compared to first movers (39) and non-adopters (17). Managerial tenure is relatively consistent across the three categories, with only minor variations in mean values. Both first movers and followers exhibit significantly higher mean values for turnover and gross value added than non-adopters, indicating that adopting firms are larger. The percentage of computer users is relatively consistent across the three categories, with minor variations in mean values.

Notably, the descriptive statistic reveals that non-adopters of cloud computing and big data differ significantly from adopters, whether they are first movers or followers, in some key characteristics: they have lower levels of investment in fixed assets, employ workers and managers with lower levels of education, have a smaller proportion of managers and workers possessing college degrees, are smaller in terms of

Table 4: Descriptive statistics – Big Data

	Non adopters		First Movers		Followers	
	Mean	P50	Mean	P50	Mean	P50
Investments in fixed assets (in ths)	81	2	257	8	515	11
Firm age	19	16	18	15	19	16
Average wage (in ths)	16	14	16	14	19	16
Capital intensity (in ths)	41	6	36	8	80	9
Education of managers	11	11	12	12	12	12
Education of workforce	10	9	10	10	11	11
Exporter (%)	7	0	10	0	10	0
Foreign-owned (%)	3	0	4	0	7	0
Leverage (%)	93	61	88	65	65	65
Previous management experience	1	0	1	0	1	0
Number of workers	17	6	39	8	63	9
Productivity (in ths)	29	19	32	20	43	22
Profitability (%)	-8	4	3	4	6	5
Share high educ. managers (%)	27	0	38	0	39	0
Share high educ. workers (%)	11	0	16	3	18	7
Tenure managers	13	12	11	10	12	11
Turnover (in ths)	2100	345	6649	556	12524	692
Gross value added (in ths)	481	106	1559	149	2715	224
Share of computer users (%)	64	67	63	67	71	83

Notes: We report all years of the sample. Sampling weights provided by the survey are used in this report. P50 stands for percentile 50. "ths" stands for thousands. *Source:* Authors' calculations based on IUTICE, QP and SCIE (INE, 2021a, INE, 2021c, and INE, 2021b).

size (measured through the number of workers, turnover, and gross value added), and show a lower inclination to foreign ownership. The differences among adopters, first movers, and followers are insignificant, but key distinctions exist between technologies. Notably, firms that adopted cloud computing and big data early tend to be slightly younger and have managers with less tenure compared to followers.

2.5 Methodology and econometric strategy

This paper aims to characterize firms adopting digital technologies and evaluate their impact on their performance. In this section, we outline the baseline models and methodology used in our estimations, taking into account the specific characteristics of our panel and the objectives of our study.

2.5.1 Characterizing cloud computing and big data adopters

To investigate the determinants of digitalization adoption, we focus on cloud computing and big data, which are defined as categorical variables. These variables take 0, 1, and 2 values to indicate the firm's

digitalization adoption behavior. A value of 1 represents early adoption, or first mover, while 2 denotes later adoption or follower, and 0 indicates non-adoption. In our analysis, we employ a Multinomial Logit Model to analyze the adoption patterns of digitalization. The probability formula for our model is presented in equation (1).

$$P(Y_{i,t} = j|X_{i,t-2}, Z_{i,t-2}) = \frac{\exp\left[X'_{i,t-2}\beta + Z'_{i,t-2}\sigma\right]}{\sum_{j=0}^{2} \exp\left[X'_{i,t-2}\beta + Z'_{i,t-2}\sigma\right]}, \quad j = 0, 1, 2. \tag{1}$$

In this equation, X represents the characteristics related to human capital, such as managers or workers, while Z denotes the control variables, including size and sector dummies, Leverage, Profitability, and Exporter Status. The index j represents the choice of adoption, which can be non-adoption, first mover or follower adopter. It's important to highlight that odds ratios in the multinomial logit model are self-contained and can be individually interpreted compared to the remaining outcomes. Our results tables present the marginal effects (at the means for continuous variables and the base category for categorical variables, as specified in the results table notes). These marginal effects reveal the difference in probability for each outcome level associated with a one-unit change in each predictor variable. Importantly, this difference is computed independently for each outcome level and remains unaffected by the base outcome specified in the multinomial logit model.

We hypothesize that various factors, including the managerial team's education level, the workforce's education level, and other relevant variables captured in a vector denoted as X, can explain the decision to adopt technology. It is important to note that our independent variables are lagged by two years. This adjustment is necessary because the IUTICE survey, which serves as the source of our dependent variable, is reported by firms at the beginning of the year. In contrast, the independent variables are derived from the QP and SCIE surveys, which reflect the firm's position at the end of the year and are collected in the subsequent year.

2.5.2 The impact of cloud computing and big data on firm performance

We used five dependent variables to analyze how digital adoption affects firm performance. These variables encompass gross value added, turnover, employment, labour productivity, and TFP, all measured using logarithmic values. It is crucial to ensure a random selection of technology adoption to establish a causal connection between technological adoption and firm performance. This requires that the groups subjected to technology adoption and those not exposed to it are nearly identical, with the only discernible difference over time being the introduction of technology within our framework, in addition to our modeled variables.

Nevertheless, we acknowledge that the treated and non-treated groups are not identical. Firms with superior performance tend to invest more in innovation, increasing their likelihood of adopting new technologies (Daveri & Parisi, 2015; Koch et al., 2021). This introduces selection issues since the two groups

under study exhibit dissimilarities. While the fixed effects specification handles selection based on time-invariant characteristics, it is vital to recognize that firm attributes can evolve due to factors unrelated to the model, thus influencing technology adoption decisions.

We adopt the approach outlined by Guadalupe et al. (2012) to address this concern and reweight firms using a propensity score estimator. This technique considers variations in the probability of adoption based on prior characteristics.

The propensity score is calculated by distinguishing between firms that adopt technology in year t as treated observations and firms that never adopt technology as control observations. Separate propensity scores are computed for each technology. To estimate the probability of firm i being a technological adopter, we pool treated and control observations from all years, accounting for exporter status, size, and capital stock (lagged by two years). The estimated probability is derived using sample survey weights, reflecting the number of firms in the sample. The resulting propensity score, denoted as p is employed to assign weights to the treated and control firms. Specifically, each treated firm is assigned a weight of $1/p$ while each control firm is assigned a weight of $1/(1-p)$.

The reweighted estimator of the propensity score is utilized in equations (2) and (4). This estimation assumes that technology adoption is random and depends on observable characteristics that may vary over time, influencing the selection process. The choice of the propensity score reweighting estimator, as opposed to the propensity score matching estimator, is justified by the findings of Busso et al. (2014), which demonstrate the superior performance of reweighting over propensity score matching in finite samples.

First movers versus non-adopters

Let $Y_{i,t}$ represent our five dependent variables for firm i in time t. Hence, the model we adopt is presented in equation (2):

$$\ln(Y)_{i,t} = \beta_1 FMAdopter_{i,t-1} + \beta_2 X_{i,t-1} + \beta_3 Z_{i,t-1} + \hat{\alpha}_i + \eta_t + \varepsilon_{i,t} \tag{2}$$

where FMAdopter represents a binary variable that takes the value of 1 if firm i is a first mover into digital technology at time $t-1$, and 0 otherwise. X is a vector of explanatory variables related to human capital characteristics, and Z is a vector of control variables (exporter, leverage, foreign-owned, importer, firm age squared, the weight of computer users, sector, and size[10]). Due to its definition, being a first mover is a constant within firms, which implies that we cannot use the standard panel data fixed effect models. In light of the extensive information captured in our datasets and following the approach proposed by Torres et al. (2018), we adopt a cross-section model with an additional control variable representing the firm-specific non-time-varying component. This component is generated by predicting the fixed effect in a

[10] In certain specifications, we exclude size and instead use intermediate inputs, capital stock, and number of employees as control variables. Please refer to the notes in the result tables for further details.

previous model. The model used for predicting the fixed effect, denoted as $\hat{\alpha}_i$, is presented in equation (3).

$$\ln(Y)_{i,t} = \beta_1 TecAdopter_{i,t-1} + \alpha_i + \eta_t + \varepsilon_{i,t} \tag{3}$$

The variable $TecAdopter$ is a binary indicator that equals 1 if firm i adopts at least one of the two new technologies covered by the IUTICE survey at time t – namely, cloud computing or big data – and 0 otherwise.

Since we have incorporated a predicted parameter, we employ the bootstrap technique to ensure valid statistical inference in the second step of the analysis (equation 2). The bootstrap application is discussed in detail in Cameron and Trivedi (2005), specifically in chapter 6.6.

First movers versus followers

To analyze the performance differences between first movers and follower adopters, we employed a Long-Differences Model to capture early technology adoption's medium-to-long-run effects. Given that all first movers in our sample adopted cloud computing in 2014 and big data in 2016, with followers entering in subsequent years, it is crucial to use a framework that mitigates year-specific fluctuations and accounts for persistent firm-level heterogeneity. This approach is particularly advantageous because, firstly, it captures the lasting effects of technology adoption, distinguishing between transient shocks and sustained performance changes. Secondly, it reduces biases from year-specific trends, ensuring that macroeconomic fluctuations do not drive observed differences. Thirdly, combined with the propensity score reweighting strategy, it controls for unobserved firm characteristics, which may influence both the timing of adoption and performance outcomes.

Previous empirical studies on technology adoption and automation have applied a long-differences estimator to assess firm-level impacts. For instance, Acemoglu et al. (2020) and Acemoglu et al. (2023) studies on robot adoption use long-differences to study firm growth and productivity shifts after technology adoption. The approach is particularly useful in distinguishing the effects of automation on firm expansion.

Our model is applied to cloud computing adoption from 2013 to 2019, with first movers being firms that adopted cloud technology in 2014 and followers reporting adoption in subsequent years. Similarly, long-differences in big data adoption were analyzed from 2015 to 2019.

Let Y_i represent our five dependent variables; our long-differences estimation equation is defined in equation (4).

$$\Delta(Y_i) = \beta_1 TecFMAdopter_i + \beta_2 X_{i,t*} + \beta_3 Z_{i,t*} + \varepsilon_i \tag{4}$$

where $\Delta(Yi)$ represents the differences in our dependent variables between the end and the beginning of the employed long-differences panel. $TecFMAdopter_i$ is a dummy variable that takes the value of

1 for first-mover firms and 0 for followers. $X_{i,t*}$ denotes a set of explanatory variables related to human capital characteristics established at the beginning of the Long-Differences panel ($t*$), $Z_{i,t*}$ represents a set of control variables (exporter, leverage, foreign-owned, importer, firm age squared, weight of computer users, sector and size) and ε_i represents the error term.

2.6 Results

In this section, we present the results of our two main research questions: (1) What characteristics distinguish first movers from followers and non-adopters of digital technologies? and (2) Is there an edge of first movers in digital technologies on firm performance?

2.6.1 Characterization of digital adopters

Table 5 presents the results for the first question of our study. The table shows the marginal effects resulting from the estimation of the multinomial logit model presented in equation (1). The table is organized with columns 1 to 3 representing cloud computing technology and columns 4 to 6 for big data technology. Marginal effects for the outcome of non-adoption (Non-Adopter (NAdopter)) in cloud computing and big data are presented in columns 1 and 4, respectively. Columns 2 and 5 contain the outcomes for first-mover adopters (First Mover (FMover)), while columns 3 and 6 are for followers. Panels A to G represent different models estimated with varying firm characteristics, analyzed separately due to multicollinearity concerns.

The first notable result is that digital non-adopters characteristics (columns 1 and 4 in Table 5) are markedly distinct from those who adopt digital technologies, whether first movers or followers. Non-adopters are more likely to have lower levels of education among managers and workers, longer managerial tenures, and represent older firms when compared to those adopting digital technologies. This finding is aligned with the ones reported in Calvino et al. (2022), Daveri and Parisi (2015), Calvino and Fontanelli (2023), among others.

Secondly, although we observe differences in the magnitude of the coefficients, the characteristics that drive firms to adopt cloud computing earlier are similar to those that prompt followers to adopt that digital technology (columns 2 and 3 in Table 5). However, when it comes to big data adoption, some distinctions emerge in the characteristics of firms at different stages of adoption. Specifically, our results suggest that firms with shorter managerial tenures and a higher proportion of college-educated workers are more likely to be pioneers in adopting big data technology compared to the remaining possible outcomes (column 5 of Table 5). Those same characteristics do not significantly impact the likelihood of becoming a follower adopter (column 6 of Table 5). However, hypothesis testing to determine differences between first movers and follower adopters of big data shows no significant statistical differences in the impact of those characteristics between these two groups.

Table 5: Selection into digitalization adoption – Multinomial logit analysis

	(1)	(2)	(3)	(4)	(5)	(6)
	Cloud Computing			Big Data		
	NAdopter	FMover	Follower	NAdopter	FMover	Follower
	Panel A					
Education manager	-0.0091***	0.0015***	0.0076***	-0.0040***	0.0024**	0.0015*
	(0.0010)	(0.0004)	(0.0009)	(0.0013)	(0.0009)	(0.0009)
	Panel B					
Share HE manager	-0.0006***	0.0001*	0.0005***	-0.0002*	0.0001	0.0001
	(0.0001)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
	Panel C					
Tenure manager	0.0017***	-0.0007***	-0.0010**	0.0017***	-0.0011**	-0.0006
	(0.0005)	(0.0002)	(0.0004)	(0.0006)	(0.0005)	(0.0004)
	Panel D					
Prev. manager exper.	-0.0035	-0.0020	0.0055**	-0.0048	0.0009	0.0039**
	(0.0026)	(0.0015)	(0.0021)	(0.0030)	(0.0024)	(0.0019)
	Panel E					
Education workers	-0.0141***	0.0025***	0.0116***	-0.0086***	0.0051***	0.0035***
	(0.0014)	(0.0006)	(0.0012)	(0.0020)	(0.0016)	(0.0012)
	Panel F					
Share HE workers	-0.0015***	0.0003***	0.0012***	-0.0009***	0.0007***	0.0002
	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0002)
	Panel G					
Young firm	-0.0223**	0.0084*	0.0139*	-0.0390***	0.0239**	0.0151*
	(0.0093)	(0.0048)	(0.0083)	(0.0129)	(0.0099)	(0.0089)
Observations	15,747			8,170		

Notes: The dependent variable in the multinomial logit model takes the value 0 if the firm is a non-adopter (NAdopter, columns 1 and 4), 1 if the firm is a technological first-mover adopter (FMover, columns 2 and 5), and 2 if the firm is a technological follower adopter (columns 3 and 6). The model includes sector and size dummies, as well as controls for profitability, leverage, and exporter status. Marginal effects are computed at mean values for continuous variables (Panels A–F) and at the base category for the dummy variable "Young Firm" (Panel G). Robust standard errors are reported in parentheses. Sampling weights from the survey are applied. *Significance levels:* *** 1%; ** 5%; * 10%. *Source:* Authors' calculations based on IUTICE, QP and SCIE (INE, 2021a, INE, 2021c, and INE, 2021b).

Looking at the results reported in Table 5, we should stress that past managerial experience (panel D of Table 5) increases the probability of being a follower in both technologies relative to non-adopting or being a first mover (a one-year increase in years of experience increases the probability of follower

adoption of cloud computing and big data by 0.6 p.p. and 0.4 p.p., respectively). This result suggests that managerial experience in previous firms may explain the higher likelihood of becoming a follower in adopting digital technologies. At the same time, deciding to be a first mover or non-adopter is inherent and potentially influenced by the firm's culture and a well-defined strategy based on previous investments. However, for big data, hypothesis testing does not allow us to reject the possibility that both first movers and followers are equally impacted by past managerial experience.

Our findings, as detailed in panels E and F of Table 5, reveal that workers' education plays a significant role in adopting digital technologies. In particular, the level of education appears to be a key determinant in becoming a first mover in big data technology and a follower in cloud computing technology.

Finally, being a young firm is particularly relevant to adopting advanced digital technologies, independently of the adoption stage (panel G of Table 5). Hypothesis testing for the equality of marginal effects indicates that being a young firm has no distinguishable impact on the likelihood of becoming either a first mover or a follower adopter of both technologies.

Focusing on first movers, the coefficients of marginal effects (columns 2 and 5 of Table 5) enables us to draw several key conclusions. A one-year increase in managers' average years of education raises the probability of first-mover adopting cloud computing and big data by 0.2 p.p., compared to the remaining possible outcomes (non-adoption or follower). An increase of one percentage point in the share of managers with higher education levels increases the probability of early adoption of cloud computing by 0.01 p.p., at a 10% significance level. However, it does not significantly impact the likelihood of adopting big data technology. Furthermore, a one-year increase in managerial tenure has a negative impact on the probability of being a first mover, reducing it by 0.07 p.p. for cloud computing and by 0.11 p.p. for big data. The effects of average years of education among workers are slightly larger than those of managers. A one-year increase in workers' education positively impacts the probability of being a first mover by 0.3 p.p. for cloud computing and 0.5 p.p. for big data. Similarly, the share of workers with a diploma follows the same probability impact trends. Being a young firm appears to affect the probability of being a first-mover in adopting big data, showing an impact of 2.4 percentage points, which is statistically significant at the 5% significance level. Being a young firm also positively impacts the probability of adopting cloud computing earlier by 0.8 p.p., though this effect is only statistically significant at the 10% level.

We find that education-related variables (Panels A, B, E, and F of Table 5) exhibit slightly stronger marginal effects for followers than for first movers in cloud computing adoption, whereas the opposite pattern emerges for big data. This counterintuitive result may reflect differences in the nature of the technologies: cloud computing is a more standardized and less knowledge-intensive technology, so early adoption may depend more on strategic orientation or external pressures than on measurable education levels, whereas followers rely more on human capital to implement the technology once it becomes widely diffused. In contrast, big data adoption is more complex and skill-intensive, which helps explain why educational endowments are more strongly associated with early adoption.

These findings provide insights that warrant further exploration of the literature on complementary intangible assets (Calvino & Fontanelli, 2023; Calvino et al., 2022), particularly regarding the role of

human capital skills. This aspect will fuel our subsequent section of results, where we will delve into the relationship between firms' performance and the potential complementary of human capital.

2.6.2 The impact of digitalization on firm performance

The second research question of our study aims to examine the impact of cloud computing and big data adoption on firm performance. Our analysis also explores potential complementarity between technology and management characteristics. Specifically, we will estimate the impact on five dependent variables: gross value added, turnover, employment, labour productivity, and TFP (all measured in logarithmic form). We will analyze how different cohorts of digital adoption affect these firm performance measures. As in our previous analysis, we will focus on three groups: first-mover adopters, followers, and non-adopters. For each digital technology, we will first examine the impact of being a first mover on firm performance compared to being a non-adopter. Subsequently, we will apply a long-differences model, similar to the approach found in the research of Acemoglu et al. (2023) and Daveri and Parisi (2015), to investigate whether there are any differential impacts of technology adoption on firm performance depending on the cohort of adoption, explicitly distinguishing between first movers and followers.

Cloud computing

The results of our econometric estimation of the impact of adopting cloud computing by first movers, presented in Table 6, show significant positive effects on firm performance. Specifically, these pioneering firms have experienced substantial improvements across various performance indicators. Prior to the inclusion of the interaction term (columns 1, 3, 5, 7, and 9 in Table 6), our findings indicate an 11% increase in Gross Value Added (GVA), a 5% boost in turnover, a 6% rise in employment, a 7% enhancement in labour productivity, and a 5% increase in TFP for first movers in comparison to non-adopters.

Furthermore, when we consider the interaction between cloud computing adoption and human capital characteristics (columns 2, 4, 6, 8, and 10 in Table 6), we do not find any evidence of a complementary relationship between the human capital characteristics under study and cloud computing on firm performance. These findings align with the results reported in Gal et al. (2019), providing consistent evidence.

When we estimate the long differences model to assess potential disparities in the impact of cloud computing adoption between first movers and followers (as depicted in Table 7), our analysis indicates that there are no discernible differences between these two groups. These results suggest that the timing of adoption did not significantly influence firms' productivity gains in the case of cloud computing technology. This observation may be attributed to the straightforward nature of cloud computing, which requires minimal expertise for implementation, incurs minimal structural costs, and possibly does not enable and require many complementary innovations. Consequently, this could explain why followers achieve a comparable impact to first movers. In this context, we highlight previous studies by Duso and Schiersch (2022) that concluded that while cloud computing improves firms' productivity, it does not impact IT investments.

Table 6: First movers versus non-adopters – Impact of Cloud Computing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	GVA		Turnover		Employment		LP		TFP	
					Panel A					
Education Manager$_{t-1}$	-0.0016	-0.0014	-0.0013	-0.0013	-0.0002	-0.0004	-0.0012	-0.0007	-0.0020	-0.0012
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Cloud Computing FM$_{t-1}$	0.1097***	0.1374**	0.0513***	0.0499	0.0609***	0.0439	0.0730***	0.1224*	0.0457***	0.1276*
	(0.014)	(0.066)	(0.012)	(0.055)	(0.010)	(0.039)	(0.011)	(0.064)	(0.012)	(0.070)
CCFM$_{t-1}$ × EM$_{t-1}$		-0.0020		0.0001		0.0013		-0.0036		-0.0061
		(0.005)		(0.004)		(0.003)		(0.004)		(0.005)
					Panel B					
Tenure Manager$_{t-1}$	0.0019***	0.0016***	-0.0009	-0.0010	0.0011***	0.0008**	0.0013**	0.0011**	0.0005	0.0005
	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)
Cloud Computing FM$_{t-1}$	0.1109***	0.0881***	0.0500***	0.0451*	0.0619***	0.0456**	0.0739***	0.0643***	0.0453***	0.0450**
	(0.014)	(0.025)	(0.012)	(0.023)	(0.009)	(0.020)	(0.011)	(0.019)	(0.011)	(0.022)
CCFM$_{t-1}$ × TM$_{t-1}$		0.0019		0.0004		0.0013		0.0008		0.0000
		(0.002)		(0.002)		(0.001)		(0.001)		(0.002)
					Panel C					
Prev. Manager Exper$_{t-1}$	-0.0074	-0.0090	0.0076**	0.0058	0.0040	0.0020	-0.0088*	-0.0094*	-0.0081*	-0.0087*
	(0.005)	(0.006)	(0.003)	(0.004)	(0.003)	(0.003)	(0.005)	(0.005)	(0.004)	(0.005)
Cloud Computing FM$_{t-1}$	0.1097***	0.1054***	0.0504***	0.0458***	0.0605***	0.0554***	0.0733***	0.0719***	0.0456***	0.0442***
	(0.013)	(0.014)	(0.012)	(0.011)	(0.010)	(0.009)	(0.011)	(0.012)	(0.011)	(0.013)
CCFM$_{t-1}$ × PME$_{t-1}$		0.0072		0.0079		0.0087		0.0024		0.0025
		(0.012)		(0.012)		(0.008)		(0.011)		(0.010)
					Panel D					
Education workers$_{t-1}$	0.0015	0.0008	-0.0001	0.0006	-0.0037***	-0.0033***	-0.0002	-0.0011	-0.0004	-0.0012
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
Cloud Computing FM$_{t-1}$	0.1085***	0.0706*	0.0509***	0.0904***	0.0625***	0.0836***	0.0728***	0.0252	0.0450***	-0.0058
	(0.014)	(0.038)	(0.012)	(0.034)	(0.010)	(0.032)	(0.011)	(0.032)	(0.011)	(0.033)
CCFM$_{t-1}$ × EW$_{t-1}$		0.0034		-0.0036		-0.0019		0.0043		0.0046
		(0.004)		(0.003)		(0.003)		(0.003)		(0.003)

Notes: All dependent variables are in logs. Labour Productivity (LP) stands for labour productivity, and TFP stands for total factor productivity. Coefficients were obtained using a Propensity Score Reweighting Estimator with sampling weights and bootstrap standard errors. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector, and computer user share. For the labour productivity, GVA, and turnover models, we additionally control for capital stock, intermediate inputs, and number of employees. For the TFP and employment models, firm size is included as a categorical control. The total number of observations is 3,366. *Significance levels:* *** 1%; ** 5%; * 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE, 2021a, INE, 2021c, INE, 2021c).

Table 7: First movers versus followers – Long-term impact of Cloud Computing

	(1)	(2)	(3)	(4)	(5)
	ΔGVA	ΔTurnover	ΔEmployment	ΔLP	ΔTFP
			Panel A		
Education manager	0.0207**	0.0249*	0.0087	0.0091	0.0116
	(0.010)	(0.013)	(0.009)	(0.008)	(0.007)
CC FM adopters	-0.0422	-0.0378	0.0119	-0.0482	-0.0424
	(0.048)	(0.045)	(0.041)	(0.036)	(0.033)
			Panel B		
Tenure manager	-0.0160***	-0.0116***	-0.0127***	-0.0028	-0.0005
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
CC FM adopters	-0.0465	-0.0416	0.0080	-0.0493	-0.0429
	(0.047)	(0.045)	(0.041)	(0.036)	(0.033)
			Panel C		
Prev. manager exper	0.0901***	0.0839***	0.0438*	0.0388**	0.0376**
	(0.025)	(0.024)	(0.024)	(0.019)	(0.017)
CC FM adopters	-0.0356	-0.0321	0.0151	-0.0453	-0.0398
	(0.048)	(0.045)	(0.041)	(0.036)	(0.033)
			Panel D		
Education workers	0.0246	0.0155	0.0292*	-0.0054	0.0027
	(0.018)	(0.016)	(0.016)	(0.011)	(0.009)
CC FM adopters	-0.0476	-0.0420	0.0081	-0.0479	-0.0431
	(0.047)	(0.045)	(0.041)	(0.036)	(0.032)

Notes: All dependent variables are in logs. LP stands for labour productivity. TFP stands for total factor productivity. Coefficients were obtained using long differences model 2013–2019 with a Propensity Score Reweighted Estimator. Sampling weights are applied with correction for variable loss. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector, and computer user share. For the labour productivity model, we additionally control for capital stock, intermediate inputs, and number of employees. For the GVA, turnover, employment, and TFP models, firm size is included as a categorical control. Robust standard errors are in parentheses. The total number of observations is 908. *Significance levels:* ***, 1%; **, 5%; *, 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE) (INE, 2021a, INE, 2021c, INE, 2021c).

Big data

Table 8 contains the results of estimating the impact of early adoption of big data on firm performance in comparison to non-adopters. We observe a consistent positive effect of early big data adoption in the odd-numbered columns across all panels (A to D). Firms that embrace big data technology experience a 10% increase in GVA, a 6% boost in turnover, a 5% rise in employment, a 6% improvement in labour productivity, and a 4% increase in TFP.

Turning to the even-numbered columns in Panel A (Table 8), where we consider the interaction between managerial characteristics and big data adoption, several key findings emerge. Firstly, there is a positive coefficient for the interaction between managers' education levels and big data adoption, particularly when explaining employment and TFP. This suggests that firms with well-educated managerial teams can reap the benefits of big data technology on TFP, and higher education levels enhance this positive impact. Specifically, a one-year increase in a manager's education level results in a 0.3 percentage point increase in the positive effects of adoption on TFP. On the other hand, we notice a slight negative effect in column 6, where the education level of the managerial team interacts with big data adoption and impacts employment. This could be attributed to improved organizational methods resulting from enhanced managerial skills, leading to increased efficiency and consequently dampening the positive effects of big data technology on employment (Bloom, Eifert, et al., 2012; Calvino & Fontanelli, 2023).

Panel B (Table 8) explores the relationship between managers' tenure and big data adoption on various performance metrics. A one-year increase in managers' tenure reduces the positive effects of big data adoption on GVA, turnover, labour productivity, and TFP by 0.20 percentage points, 0.35 percentage points, 0.14 percentage points, and 0.15 percentage points, respectively. It's worth noting that the positive impact of big data technology will persist in firms with the most tenured managers; it will only be reduced.

Moving on to Panel C (Table 8), we observe a positive interaction between past managerial experience and big data adoption on turnover. Firms benefit from the combination of experienced managers and the utilization of big data, resulting in an increased positive impact of early adoption on turnover by 0.68 percentage points.

Finally, in Panel D (Table 8), we consider the education levels of workers. Similarly to managerial education, we find a significant positive coefficient for the interaction term in models for firms' productivity in columns 2, 8, and 10. This indicates that firms can only harness the productivity gains associated with big data technology if they have a workforce with a certain level of education. This positive impact is observed in GVA, labour productivity, and TFP by 0.39 percentage points, 0.42 percentage points, and 0.39 percentage points, respectively, with the first and last being statistically significant at the 10% significance level.

In summary, the findings in Table 8 emphasize the importance of managerial and workforce characteristics and their interactions with big data adoption in shaping firm performance outcomes. These results align with the expectations outlined in existing literature, including studies of Niebel et al. (2019), Schivardi and Schmitz (2019), Daveri and Parisi (2015), and others, providing further support for the notion

that firms with higher human capital levels are more likely to benefit from technological implementations, leading to improved performance outcomes.

When examining the long differences model that compares the performance of first movers in big data technology with followers, presented in Table 9, our analysis uncovers significant differences. Firms that embraced big data technology early on (first movers) experienced substantially more positive impacts on GVA, LP, and TFP than those that adopted it later. Specifically, these pioneering firms witnessed notable improvements in GVA, LP, and TFP by approximately 10%, 9%, and 6%, respectively. These findings highlight the edge of being a first mover in adopting big data technology, as it leads to significantly enhanced performance outcomes.

Interestingly, and consistent with the evidence reported in Table 5, these performance results suggest that, in the case of cloud computing, followers rely more on measurable human capital endowments when adopting the technology. However, once adoption takes place, productivity gains are comparable between first movers and followers, reinforcing the idea that timing plays a limited role for this relatively standardized technology. Moreover, the absence of significant complementarities between human capital variables and cloud computing adoption in firm performance (as reported in columns 2, 4, 6, 8, and 10 of Table 6) supports this interpretation.

In contrast, our findings for big data indicate that education and other human capital characteristics play a complementary role in generating productivity gains (as reported in columns 2, 4, 6, 8, and 10 of Table 8), consistent with the stronger association between these variables and first-mover adoption observed in Table 5. Furthermore, for big data, we find that first movers achieve significantly higher productivity gains than followers (Appendix Table 39), in sharp contrast with the pattern observed for cloud computing. This result highlights that the importance of adoption timing depends on the characteristics of the technology and on the complementarities required to unlock its full potential.

These results align with those found in Niebel et al. (2019). It is interesting to note that, unlike cloud computing technology, there is a different payoff for being an early adopter compared to being a follower in the case of big data. This difference in outcome could be attributed to the greater complexity of big data technology for implementation and utilization, which requires specific skills among human capital to realize its potential benefits. In this scenario, first movers have a timing advantage as they utilize this new technology sooner. Furthermore, considering the findings of Niebel et al. (2019) indicating that big data enhances innovation levels within firms, early adopters are expected to accelerate the intensity of other innovation sources sooner than followers. This emphasizes the complementary nature of the intangible assets facilitated by big data technologies, a factor that may be reflected in these results.

As part of the robustness checks, we present in the appendix (Tables 36, 37, 38, and 39) the results without propensity score reweighting. This allows us to test the sensitivity of our findings. While we acknowledge the dissimilarities between the groups, the robustness checks show that omitting propensity score reweighting does not significantly affect the main conclusions, indicating that the findings remain

Table 8: First movers versus non-adopters – Impact of Big Data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	GVA		Turnover		Employment		LP		TFP	
Panel A										
Education Manager$_{t-1}$	0.0013*	0.0011	-0.0006	-0.0007	0.0003	0.0007	0.0013*	0.0010	0.0015*	0.0009
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Big Data FM$_{t-1}$	0.0975***	0.0803***	0.0631***	0.0584***	0.0456***	0.0776***	0.0614***	0.0377*	0.0389***	-0.0052
	(0.008)	(0.025)	(0.007)	(0.022)	(0.005)	(0.014)	(0.006)	(0.019)	(0.006)	(0.019)
BD$_{t-1}$ × EM$_{t-1}$		0.0013		0.0004		-0.0024**		0.0018		0.0033*
		(0.002)		(0.002)		(0.001)		(0.002)		(0.002)
Panel B										
Tenure Manager$_{t-1}$	-0.0002	0.0002	-0.0006*	0.0001	0.0001	0.0001	-0.0001	0.0001	-0.0003	0.0000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Big Data FM$_{t-1}$	0.0979***	0.1229***	0.0616***	0.1053***	0.0458***	0.0516***	0.0618***	0.0794***	0.0393***	0.0576***
	(0.008)	(0.014)	(0.008)	(0.017)	(0.005)	(0.011)	(0.006)	(0.009)	(0.006)	(0.008)
BD$_{t-1}$ × TM$_{t-1}$		-0.0020**		-0.0035***		-0.0005		-0.0014**		-0.0015***
		(0.001)		(0.001)		(0.001)		(0.001)		(0.001)
Panel C										
Prev. Manager Exper$_{t-1}$	0.0008	-0.0005	0.0011	-0.0009	0.0023**	0.0018	-0.0000	-0.0009	-0.0006	-0.0017
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
Big Data FM$_{t-1}$	0.0980***	0.0943***	0.0622***	0.0567***	0.0453***	0.0440***	0.0620***	0.0597***	0.0397***	0.0366***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)
BD$_{t-1}$ × PME$_{t-1}$		0.0045		0.0068**		0.0016		0.0029		0.0039
		(0.004)		(0.003)		(0.002)		(0.004)		(0.004)
Panel D										
Educ workers$_{t-1}$	0.0033**	0.0025	0.0005	0.0005	0.0012	0.0015**	0.0028**	0.0019	0.0012	0.0004
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Big Data FM$_{t-1}$	0.0973***	0.0556**	0.0623***	0.0615***	0.0454***	0.0591***	0.0614***	0.0165	0.0393***	-0.0024
	(0.008)	(0.022)	(0.008)	(0.022)	(0.005)	(0.014)	(0.006)	(0.020)	(0.006)	(0.020)
BD$_{t-1}$ × EW$_{t-1}$		0.0039*		0.0001		-0.0013		0.0042**		0.0039*
		(0.002)		(0.002)		(0.002)		(0.002)		(0.002)

Notes: All dependent variables are in logs. LP stands for labour productivity, and TFP stands for total factor productivity. Coefficients were obtained using a Propensity Score Reweighting Estimator with sampling weights and bootstrap standard errors. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector, and computer user share. For the labour productivity, GVA, and turnover models, we additionally control for capital stock, intermediate inputs, and number of employees. For the TFP and employment models, firm size is included as a categorical control. The total number of observations is 3,891. *Significance levels:* *** 1%; ** 5%; * 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE, 2021a, INE, 2021c, INE, 2021c).

Table 9: First movers versus followers – Long-term impact of Big Data

	(1)	(2)	(3)	(4)	(5)
	ΔGVA	ΔTurnover	ΔEmployment	ΔLP	ΔTFP
			Panel A		
Education manager	-0.0008	0.0154	0.0098	-0.0116	-0.0093
	(0.010)	(0.011)	(0.007)	(0.007)	(0.007)
Big data FM adopters	0.1069**	0.0347	0.0211	0.0864**	0.0549*
	(0.045)	(0.046)	(0.036)	(0.041)	(0.033)
			Panel B		
Tenure manager	-0.0085***	-0.0076***	-0.0071***	-0.0012	-0.0003
	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
Big data FM adopters	0.0999**	0.0296	0.0145	0.0845**	0.0544
	(0.045)	(0.047)	(0.036)	(0.041)	(0.033)
			Panel C		
Prev. manager exper	0.0264	0.0278	0.0143	0.0136	0.0028
	(0.019)	(0.019)	(0.014)	(0.012)	(0.009)
Big data FM adopters	0.1049**	0.0338	0.0208	0.0845**	0.0546*
	(0.045)	(0.047)	(0.037)	(0.041)	(0.033)
			Panel D		
Education workers	0.0331**	0.0127	0.0182	0.0160	0.0116
	(0.016)	(0.014)	(0.011)	(0.011)	(0.010)
Big data FM adopters	0.1041**	0.0348	0.0214	0.0842**	0.0549*
	(0.044)	(0.046)	(0.036)	(0.041)	(0.033)

Notes: All dependent variables are in logs. LP stands for labour productivity. TFP stands for total factor productivity. Coefficients were obtained using long differences model 2013–2019 with Propensity Score Reweighted Estimator. Sampling weights are applied with correction for variable loss. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector, and computer user share. For the labour productivity model, we additionally control for capital stock, intermediate inputs, and number of employees. For the GVA, turnover, employment, and TFP models, firm size is included as a categorical control. Robust standard errors are in parentheses. The total number of observations is 773. *Significance levels:* ***, 1%; **, 5%; *, 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE, 2021a, INE, 2021c, INE, 2021c).

consistent[11].

2.7 Concluding remarks

This paper has examined which firms adopt cloud computing and big data analytics, when they adopt them, and how adoption timing translates into economic performance. Using matched Portuguese firm-level data on ICT use, balance-sheet outcomes, and linked employer–employee records, we classified adopters as first movers or followers and compared them to non-adopters.

Our results indicate that adoption patterns and their effects unfold through two complementary mechanisms. First, internal capabilities play a decisive role in both adoption and its payoffs (supporting Hypothesis 1, formulated in Section 2.3). Firms adopting either technology tend to be younger and larger, employing more highly educated managers and workers. Managerial education is positively associated with the likelihood of adopting cloud computing and big data analytics, whereas longer tenure slows adoption (consistent with Hypothesis 2). Prior managerial experience increases the probability of being a follower rather than a pioneer, suggesting a more cautious, long-term approach to ICT investment. These findings highlight the importance of intangible assets and human capital – particularly leadership skills and adaptive capacity – in shaping firms' digital strategies and enabling them to translate adoption into superior outcomes.

Second, organizational learning generates cumulative advantages for early movers, particularly for more complex technologies. While digitalization yields substantial performance gains overall, the size and persistence of these gains differ across technologies. For cloud computing, first movers experience improvements in gross value added, turnover, employment, and productivity relative to non-adopters, but followers quickly catch up – consistent with the view that cloud computing is a relatively modular and accessible technology with limited learning requirements (supporting Hypothesis 4). In contrast, big data analytics exhibits a more capability-dependent pattern: productivity gains from early adoption are stronger when firms have well-educated managers and workers (supporting Hypothesis 3), and, over time, first movers outperform followers in both value added and productivity (again supporting Hypothesis 4). This indicates that big data adoption involves a significant learning curve and firm-specific knowledge accumulation that are difficult for followers to replicate.

This study makes several contributions to the existing literature on digital technology adoption. Unlike prior research, which typically focuses on either the determinants of adoption or the productivity effects of cloud computing and big data in isolation (e.g., Bloom, Eifert, et al., 2012; Brynjolfsson, Jin, & McElheran, 2021; Calvino et al., 2022), we integrate both dimensions and explicitly distinguish between first movers, followers, and non-adopters. By combining detailed firm-level data on managerial and workforce

[11] As a robustness check, we re-estimated all main specifications after excluding (i) the manufacturing sector and (ii) the wholesale and retail sector – the two largest sectors in our sample. The main conclusions remain unchanged: the systematic differences between adopters and non-adopters persist, cloud computing shows no first-mover advantage, and big data exhibits capability-dependent returns with advantages for early adopters. Detailed results are available upon request.

characteristics with information on adoption timing, we provide new evidence on how internal capabilities and organizational learning jointly shape the returns to digital technologies. This approach allows us to uncover not only which firms adopt early and why, but also how adoption timing interacts with human capital to generate sustained performance advantages – particularly for complex technologies such as big data analytics. Taken together, our findings support a theoretical framing in which internal capabilities and organizational learning jointly determine the adoption and effective use of digital technologies, offering a lens through which to interpret heterogeneity in adoption timing and performance outcomes.

These findings carry implications for both policymakers and managers. For policymakers, the results highlight two priorities: (i) expanding access to advanced managerial and analytical skills – through education, training, and targeted support – can help close the gap between digital leaders and laggards; and (ii) improving the quality and scope of firm-level data on technology use would support more informed policy design and enable closer monitoring of the economic and labour market impacts of digitalization. For managers, the evidence suggests that early adoption of low-complexity technologies such as cloud computing offers measurable but temporary benefits, as followers can quickly replicate these gains. In contrast, early adoption of more complex technologies such as big data yields more persistent performance improvements – particularly when supported by a skilled and adaptable leadership team – highlighting the importance of aligning technology investment with internal capabilities and fostering organizational learning.

Finally, two limitations should be acknowledged. First, the analysis is restricted to two technologies within a single national context, and its findings should not be generalized without caution. Second, the temporal structure of the data limits our ability to capture long-run dynamics and transitions. Future research should leverage richer and more frequent data to track the diffusion of emerging technologies – such as artificial intelligence – and to examine how the interplay between multiple digital tools influences firm performance. Understanding the long-term consequences of adoption timing and the cumulative nature of learning remains a central priority for advancing this research agenda.

3

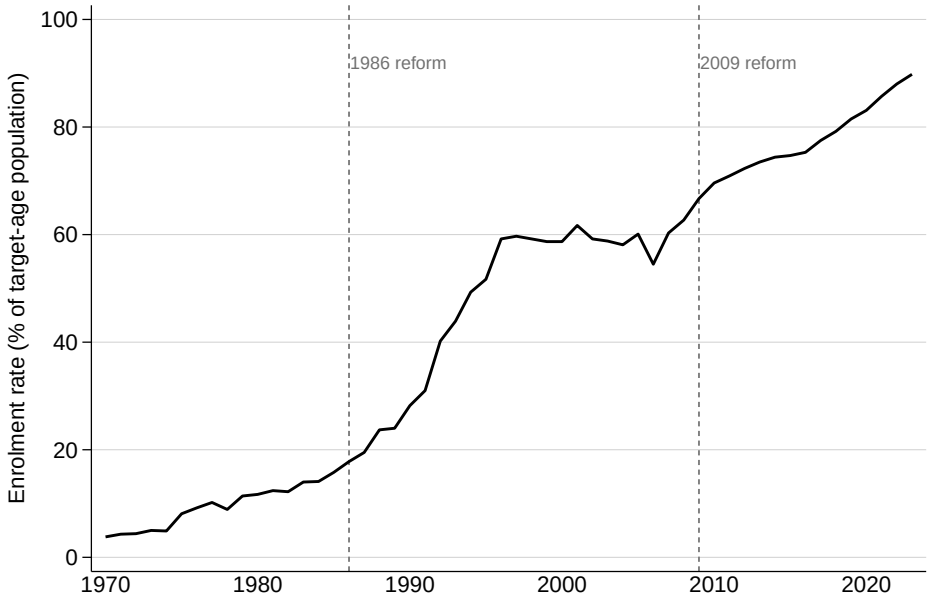
Keeping teenagers in school and firm adjustment: evidence from Portuguese manufacturing firms

3.1 Introduction

How do compulsory schooling reforms affect firms? The private and social returns to education have been extensively studied, but less is known about how these reforms change firms' production decisions by restricting the supply of young, low-skilled labour.[1] In this chapter, we study Portugal's 2009 reform – which extended mandatory education from 9 to 12 years and raised the minimum school-leaving age from 15 to 18 – to examine whether the withdrawal of underage workers from the labour market is associated with differential firm adjustment in manufacturing.

We examine three distinct margins of firm adjustment: robot importing through international trade records, internal occupational recomposition measured through employer–employee occupation data, and complementary ICT employment. These outcomes are not interchangeable. Robot imports capture equipment acquisition. The occupation proxy captures whether the firm employs workers in automation-related occupations. ICT employment captures complementary technical labour demand and does not directly measure robot adoption. The mechanism linking the reform to these margins is a labour-supply restriction: as minors are required to remain in school, firms that relied on underage workers lose access to a specific, low-cost input. The directed technical change literature suggests that such factor scarcity may redirect technology choices (Acemoglu, 2002, 2010; Danzer et al., 2024; San, 2023), but the shock here is narrow – minors accounted for roughly 3–5% of the workforce in firms that employed at least one worker under age 18 before the reform (hereafter, "treated" firms) – so the scale of any downstream response is an empirical question. In the firm fixed-effects specifications, complementary ICT employment is the clearest downstream margin: more exposed firms are 1.2–1.6 percentage points more likely to employ ICT workers after the reform. An occupation-based proxy is also positive, though less precisely estimated. Annual robot-import indicators remain close to zero. We also examine firm exit as a secondary margin.

[1] This chapter is co-authored with Fernando Alexandre and Miguel Portela.



Notes: The figure plots the secondary-school real-enrolment-rate series relative to the target-age population. Year t denotes the end year of school year $t-1/t$; for example, 2023 corresponds to 2022/23. *Source:* Official DGEEC statistical publications (DGEEC, 2023, 2024).

Figure 1: Enrolment rates relative to target-age population, 1970–2023 (%)

Portugal provides a suitable setting. The 2009 reform was the most recent in a sequence of legislative changes that have raised mandatory schooling since the mid-twentieth century (Table 10). Figure 1 shows large increases in enrolment rates after both the 1986 and 2009 reforms, consistent with the latter being binding for the affected cohorts – teenagers who, under the previous regime, could have left school at age 15 and entered the labour market.

Table 10: Mandatory schooling reforms in Portugal

Law	Min educ	Min age	Children born	Entry school	LM shortage
42994/60	4		Until 1966*		
45810/64	6	14	Between 1967 and 1980**	1973/1974	1979
46/86	9	15	Between 1981 and 1996***	1987/1988	1993
85/2009	12	18	From 1997 onwards	2003/2004	2012

Notes: LM stands for Labour Market. The entries on minimum compulsory schooling and minimum school-leaving age are compiled by the authors from Laws 42994/60, 45810/64, 46/86, and 85/2009. The birth-cohort mapping, school-entry timing, and labour-market-shortage timing are authors' calculations based on the legislation's scope and transition rules. * Entry at age 7 until December 31st; **entry at age 7 until December 31st for cohorts born up to 1973 and at age 6 for later cohorts; *** entry at age 6 until September 30th.
Source: Authors' compilation from Laws 42994/60, 45810/64, 46/86, and 85/2009.

This reform is relevant for manufacturing, which has historically relied on underage labour – especially in northern Portugal, where lower educational attainment and a strong manufacturing base led minors to enter the labour market early (Goulart & Bedi, 2017). Section 3.3 shows a large post-reform shift in the

composition of new manufacturing entrants, consistent with the reform having tightened access to young, low-skilled labour.

To study firms' responses, we use a matched employer–employee dataset (*Quadros de Pessoal*) linked through an anonymized firm identifier to firm-level international trade records (*Comércio Internacional*). The trade data allow us to measure industrial robot imports at the firm level, following Acemoglu and Restrepo (2018a) and Acemoglu et al. (2023). The employer–employee records also provide the occupation-based and ICT indicators used as additional outcomes. These capture distinct internal workforce margins rather than alternative direct measures of the same underlying technology. The regression panel covers 2002–2019, avoiding the COVID-19 pandemic and placing some distance from the earlier 1986 compulsory-schooling reform (whose labour-market effects became binding around 1993). The descriptive evidence uses the widest window allowed by each source, including entrant-composition series back to 1991.

We adopt a common-timing Difference-in-Differences (DiD) strategy that compares post-reform changes between firms that were more versus less exposed to underage labour before the reform, measured using firms' employment of minors during 2006–2008. We use two post-period indicators: Post-2009, which captures enactment and possible anticipation, and Post-2012, which captures the later phase in which the labour-supply restriction should have become more directly binding. The resulting coefficients are relative reduced-form effects, not effects against a fully untreated control group. The sector–region cell analysis is treated only as complementary aggregate evidence.

The chapter contributes to three literatures. First, it extends work on compulsory schooling reforms – which has focused on educational attainment, earnings, and child labour – to the firm side (Angrist & Krueger, 1991; Dayıoğlu & Kırdar, 2022; Oreopoulos, 2006). Second, it connects to the firm-level automation literature, where the question of who adopts robots and what adoption implies for labour demand is well established (Acemoglu et al., 2023; Koch et al., 2021). Third, it relates to directed technical change, in which factor scarcity redirects technology choices (Acemoglu, 2010; Acemoglu & Restrepo, 2022a; Carneiro et al., 2023; San, 2023). Our setting combines a well-documented first stage with a narrow labour-supply shock. The main downstream pattern is higher complementary ICT employment among more exposed firms. The occupation-based proxy is also positive but less precisely estimated, and the import-based robot margins remain null. Interpretation is bounded by three first-order concerns. More- and less-exposed firms are not perfectly comparable even within a fixed-effects design. The reform also overlaps with Portugal's sovereign-debt crisis. In addition, the ICT result is not reform-specific in a clean sense: age-based placebo exposures produce ICT coefficients around 0.023–0.024, larger than the baseline 0.012–0.016, fake-timing exercises yield positive pseudo-effects in some pre-reform years, and trend-augmented TWFE and matched designs attenuate the coefficient sharply. The chapter therefore provides evidence on differential post-reform firm adjustment rather than clean evidence of robot adoption.

As a thesis chapter, it establishes the firm-side empirical question, the linked-data architecture, and the main reduced-form patterns while keeping the design's main vulnerabilities in view.

The remainder of this chapter is organised as follows. Section 3.2 reviews the related literature. Section 3.3 documents the institutional context and descriptive evidence on education and labour market entry. Section 3.4 describes the data sources and variable construction. Section 3.5 sets out the empirical strategy. Section 3.6 presents the main results, and Section 3.7 concludes.

3.2 Related literature

This chapter draws on three literatures: compulsory schooling reforms, firm-level automation, and directed technical change under factor scarcity. The question linking them is whether a schooling reform that keeps minors out of the labour market also changes firms' technology choices. Below we review each strand, focusing on the results and methods most relevant to our design.

3.2.1 Compulsory schooling reforms and labour supply

The classic literature uses compulsory schooling laws to identify the effects of education on later outcomes. Angrist and Krueger (1991) use quarter-of-birth variation to show that compulsory-attendance laws keep some potential dropouts in school and yield instrumental-variables estimates of the return to education close to OLS estimates. Oreopoulos (2006) finds that when compulsory schooling laws bind more sharply, the effects on schooling and adult earnings can be large, though the magnitude varies across settings and cohorts (Devereux & Hart, 2010). Pischke and von Wachter (2008) find zero wage returns to compulsory schooling in Germany, suggesting that the productive content of additional years depends on the institutional context. More recent papers extend the focus beyond average earnings. In Turkey, compulsory schooling reduced child labour and increased school attendance, with stronger effects in rural areas and among households with lower parental education (Dayıoğlu & Kırdar, 2022). In Mexico, the 1993 reform raised schooling but had limited effects on earnings and employment (Leon-Bravo, 2022). In Poland, prolonging general education shifted students toward general tracks; wage effects were mixed, with gains for rural vocational graduates but negative effects for secondary general graduates (Strawinski & Broniatowska, 2024).

This literature shows that schooling mandates change both the timing and the composition of labour supply, but it has rarely examined how firms respond to the withdrawal of young workers. From a firm's perspective, a binding increase in the school-leaving age is a negative supply shock for a specific labour input – young, low-cost workers – rather than a human-capital improvement for returning workers. This Chapter studies this firm-side margin directly.

The Portuguese context is well suited because the literature documents both high returns to schooling and a historically important role for early labour market entry. Vieira (1999) and Campos and Reis (2017) report high private returns to education in Portugal, and Goulart and Bedi (2017) shows that child labour remained economically important well into the late twentieth century.

3.2.2 Automation, robots, and firm behaviour

A growing literature studies what robots do inside firms and which firms adopt them. The task-based framework of Autor et al. (2003) and Acemoglu and Autor (2011) shows that automation replaces workers in routine tasks while raising the relative importance of complementary, non-routine skills. Acemoglu and Restrepo (2022b) estimate that 50–70% of the changes in the US wage structure over four decades are accounted for by automation-driven task displacement. Autor (2019) documents long-run occupational polarization and shows that among non-college workers the erosion of middle-skill jobs has pushed employment toward lower-paid service work. At the aggregate level, Acemoglu and Restrepo (2020) find that robot exposure reduces employment and wages in more exposed US labour markets, while Graetz and Michaels (2018) estimate that industrial robots raised labour productivity by about 0.36 percentage points per year across 17 countries without significantly reducing total employment. Dauth et al. (2021) report a similar pattern in Germany: robot-driven displacement in manufacturing was offset by job creation in services, with the adjustment falling mainly on younger cohorts entering the labour market.

For this chapter, the firm-level evidence is most directly relevant. Using Spanish manufacturing data, Koch et al. (2021) show that robot adopters are positively selected on output and labour productivity, and that adoption increases output by 20–25% while lowering the labour cost share by 5–7 percentage points. Using Dutch matched employer-employee data, Acemoglu et al. (2023) find positive effects for adopting firms but negative spillovers on competitors. Bonfiglioli et al. (2024) use French firm-level robot import data and show that, once demand-driven selection is addressed, exogenous exposure to automation leads to job losses at adopting firms. Humlum (2022) estimates a structural model using Danish data and finds that robot adoption expands output by 20% but shifts labour demand from production to technical workers, suggesting that within-firm occupational restructuring is itself an adjustment margin alongside equipment acquisition. Benmelech and Zator (2025) show that robot adoption in German firms responds to labour scarcity, measured through firms' self-reported difficulty in finding workers. These papers show that automation is not randomly assigned across firms and is not neutral for labour demand. Adoption is concentrated among firms that can reorganise production and employ complementary skills, which is why a policy-induced restriction on the supply of young, low-wage labour may affect robot adoption, occupation-based workforce reorganisation, and complementary ICT employment.

A measurement issue that is important for our design concerns how firm-level automation is observed. Following Acemoglu and Restrepo (2018a) and Acemoglu et al. (2023), we identify robot-importing firms through commodity code 847950 in trade records. Bonfiglioli et al. (2024) and Leone (2024) use similar trade-based measures in France and cross-country settings. Because import data capture only equipment acquisition and are extremely sparse at the firm level, we complement them with occupation-based indicators from matched employer-employee records that measure whether a firm employs workers in automation-related occupations. These two layers capture different margins – equipment purchase versus internal workforce composition – and are only weakly correlated. Chen and Roth (2024) document limitations of $\log(1+Y)$ transformations when the outcome has a mass at zero, which is relevant to our

cumulative import measures.

3.2.3 Labour scarcity and induced technological change

The mechanism we study is closest to the literature on directed technical change. In Acemoglu (1998), Acemoglu (2002), and Acemoglu (2010), the direction of innovation is endogenous to relative factor supplies. Acemoglu (2002) formalises two competing forces: a price effect, which favours innovations directed at scarce factors, and a market-size effect, which favours technologies used by abundant factors. When the elasticity of substitution is sufficiently high, the price effect dominates, so that labour scarcity encourages labour-saving innovation.

Recent empirical evidence supports this prediction across different settings. Acemoglu and Restrepo (2022a) show that ageing economies adopt more robots in places and industries more exposed to shortages of middle-aged workers; demographic change alone accounts for about 35% of the cross-country variation in robot adoption and close to half when other automation technologies are included. San (2023) finds that the termination of the *Bracero* programme in the United States increased innovation in technologies related to the most exposed crops, with effects persisting more than 15 years after the shock. M. A. Clemens et al. (2018) study the same episode and find that the exclusion of Mexican bracero workers prompted adoption of less labour-intensive technologies and shifts in crop mix, with limited effects on domestic wages. Mann and Pozzoli (2024) show that a higher share of non-Western immigrants in Danish municipalities reduces firm-level robot adoption, consistent with immigration and automation acting as substitutes. Danzer et al. (2024) exploit a German immigrant allocation policy and find that an increase in the low-skilled workforce *reduces* automation innovation (measured by patents), with the largest effects in industries employing many low-skilled workers and in tight labour markets – the mirror image of the scarcity-induced adoption we study.

Labour costs can also shape the automation margin through institutional channels. Lordan and Neumark (2018) find that minimum wage increases reduce the share of automatable employment among low-skilled workers, and Hershbein and Kahn (2018) find that recessions accelerate routine-biased technological change, with firms raising skill requirements and increasing capital investment in harder-hit areas. In a related setting, Carneiro et al. (2023) find that an exogenous increase in the supply of college-educated workers in Norway raised the productivity of skilled labour and increased R&D investment, consistent with firms adopting skill-complementary technologies in response to a positive skill-supply shock.

The common finding in this literature is that changes in factor supply – whether driven by demography, migration policy, education reform, or macroeconomic shocks – can change firms' technology choices. This chapter applies the same logic to a different setting. Rather than studying ageing, migration, or college expansion, we examine whether a compulsory schooling reform that delayed the labour market entry of minors changed firms' automation choices. The shock is distinctive because it removes a *specific, low-cost labour input* – underage workers – rather than creating a general labour shortage. The chapter therefore contributes evidence on the firm-side consequences of schooling policy: its potential effect on

technology adoption and complementary skill demand, with firm exit as a secondary margin.

3.3 Institutional context

This section describes the institutional background and presents descriptive evidence consistent with the reform having tightened the supply of young labour. We first describe the 2009 reform in its historical context, including the concurrent macroeconomic developments that any empirical analysis of this period must address. We then document two patterns: a large increase in high-school enrolment across all regions, and a shift in the composition of new manufacturing entrants toward higher education. Together, these patterns motivate the labour-supply-shock interpretation used in the firm-level analysis.

3.3.1 The 2009 reform in context

Portugal's educational attainment has historically lagged behind its European peers, a legacy of late democratisation and decades of underinvestment in schooling during the Estado Novo period. Compulsory education was limited to 4 years until 1964, when it was extended to 6 years. The landmark *Lei de Bases do Sistema Educativo* (Law 46/86), enacted two years after EU accession, established 9 years of compulsory schooling – still below the European norm. By the late 2000s, around 72% of Portuguese adults aged 25–64 had not completed upper secondary education, compared to an OECD average of 29% (OECD, 2010).

Against this backdrop, Portugal's early school-leaving rate – the share of 18 –24 year-olds with at most lower secondary education who are not in further education or training – stood at approximately 35% in 2008, more than double the EU-27 average of 14.4% (Eurostat, 2024). The gap was especially wide for young men and in regions where manufacturing employment offered a viable alternative to continued schooling.

Law 85/2009, enacted on 27 August 2009, was the largest extension of compulsory education in Portuguese history. It raised the mandatory schooling requirement from 9 to 12 years and increased the minimum school-leaving age from 15 to 18, also establishing the universality of pre-school education from age 5 (Table 10). With this reform, Portugal joined the small group of European education systems mandating at least 12 years of full-time compulsory education. Children born from 1997 onwards were the first fully affected cohort; they entered primary school in 2003/2004 and would have reached the previous school-leaving age of 15 around 2012, at which point the new requirement to stay in school until 18 became binding. The reform's effect on early school leaving was large: Portugal's rate fell from 28.3% in 2010 to 13.5% by 2015 and reached 6.4% in 2021, remaining below the EU average by that point (Eurostat, 2024). Figure 1 places this reform in longer historical perspective: secondary-school enrolment relative to the target-age population rises after both the 1986 and 2009 reforms, consistent with a sequence of increasingly binding schooling requirements.

The manufacturing sector is a natural setting to study this reform's firm-side effects. Although manufacturing's share of GDP had declined to approximately 11% by 2009, the sector remained a major employer, particularly in the North, where the Porto–Braga–Aveiro industrial corridor concentrates much of Portugal's capacity in textiles, footwear, metalworking, and automotive components. Historically, these labour-intensive industries absorbed young workers with low levels of schooling, and minor labour remained economically significant in northern Portugal well into the 2000s. Goulart and Bedi (2017) document that the North accounted for over half of all minor labour in Portugal as late as 2001, with the textiles and footwear sectors among the largest employers of undocumented youth. The 2009 reform therefore interacted with a combination of low schooling, early labour market entry, and manufacturing dependence that was especially present in this region.

3.3.2 Enrolment response

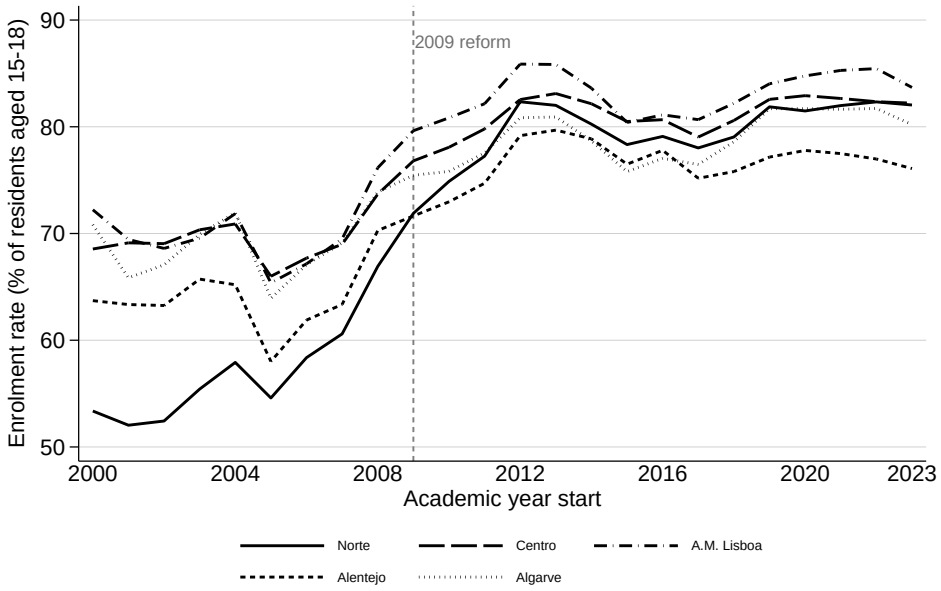
Figure 2 plots the share of high-school enrolment relative to the target population (residents aged 15–18) by Nomenclature of Territorial Units for Statistics (NUTS 2) region. Prior to the reform, enrolment rates were well below full coverage, with the North below the other regions over most of the pre-reform period. Over the post-2009 period, enrolment levels are higher in all regions than at the start of the sample, with the North showing the largest gain: in the reconstituted series used here, it moves from roughly 53% in 2000 to about 82% by 2023. Lisbon and the southern regions also rise, but more gradually. This pattern is consistent with a regime in which the reform bound more strongly where early labour-market entry had historically been more common.

The response in the North is important for our analysis because this region is one of the most manufacturing-intensive in Portugal and one of the places where early labour-market entry and the employment of minors were most common.

Two concurrent institutional developments should be noted. First, the *Novas Oportunidades* programme, launched in 2005 as a joint initiative of the Ministries of Education and Labour, expanded access to adult education through a network of approximately 450 recognition centres that eventually registered 1.6 million enrolments and issued over 400,000 certifications. This programme may have contributed to broader enrolment trends, but it targeted adults, not the school-age population. Second, Portugal's sovereign debt crisis (2009–2013) – which culminated in a €78 billion bailout programme in May 2011 and drove youth unemployment to 38% by 2013 (Reis, 2013) – raised the opportunity cost of leaving school early and may have independently increased enrolment. To focus on the school-age population most directly targeted by the reform, the enrolment metrics in Figure 2 exclude adult-specific programmes.

Table 11 provides a simple descriptive regression summary of regional post-reform shifts in enrolment at the NUTS 2 level. This table is not part of the firm-level identification strategy; its purpose is only to formalise the pattern shown in Figure 2. Under the post-2009 coding, the largest estimated shifts are in the North (0.26 log points) and Lisbon (0.19 log points). Under the post-2012 coding, Lisbon is slightly larger than the North, but the same two regions remain at the top. The broad message is therefore

consistent with the figure: the rise in high-school enrolment was widespread, but particularly pronounced in the regions where early labour-market entry had historically been more salient.



Notes: The figure plots youth-oriented upper-secondary enrolments divided by the resident population aged 15–18 in each region. Academic years are indexed by their start year on the x-axis. *Sources:* DGEEC upper-secondary enrolment tables and archive workbooks (Direção-Geral de Estatísticas da Educação e Ciência (DGEEC), 2026) and INE resident-population indicator 0000256 (Instituto Nacional de Estatística (INE), 2026).

Figure 2: High-school enrolment rate relative to resident population aged 15–18 by NUTS 2 region, 2000–2023 (%)

3.3.3 Young workers in manufacturing

We next examine how the composition of new entrants into manufacturing employment evolved over the period 1991–2021, where "new entrants" are workers entering the Portuguese private labour market for the first time. This evidence is relevant for the interpretation of the firm-level design because it shows how the pool of potential entrants into manufacturing shifted toward higher education over time.

Figure 3 illustrates the educational progression of new manufacturing entrants. In 1991, workers with up to a 9th-grade education represented 90% of all entries, while high-school graduates and tertiary-educated workers accounted for 7% and 3%, respectively. By 2009, the shares had shifted to 63%, 24%, and 13%. The long-run trend toward more educated entrants continued thereafter, reaching 35% up to 9th grade, 47% high school, and 18% tertiary by 2021. Year-to-year movements around the early 2010s should not be over-read relative to this broader pattern. Overall, the figure points to a substantial long-run transformation of the skill profile of manufacturing's labour supply.

Table 11: Regional post-reform shifts in high-school enrolment (NUTS 2)

	Post-period definition	
	(1) 2009 onward	(2) 2012 onward
Post × Norte	0.2576***	0.1567***
	(0.036)	(0.048)
Post × Algarve	0.1410***	0.1166***
	(0.027)	(0.029)
Post × Centro	0.1308***	0.0701**
	(0.021)	(0.027)
Post × Lisbon	0.1874***	0.1661***
	(0.028)	(0.033)
Post × Alentejo	0.1612***	0.1043***
	(0.032)	(0.030)
Log population aged 15–18	0.7384***	0.4711***
	(0.090)	(0.150)
Observations	100	100
R-squared	0.998	0.996
Region FE	Yes	Yes

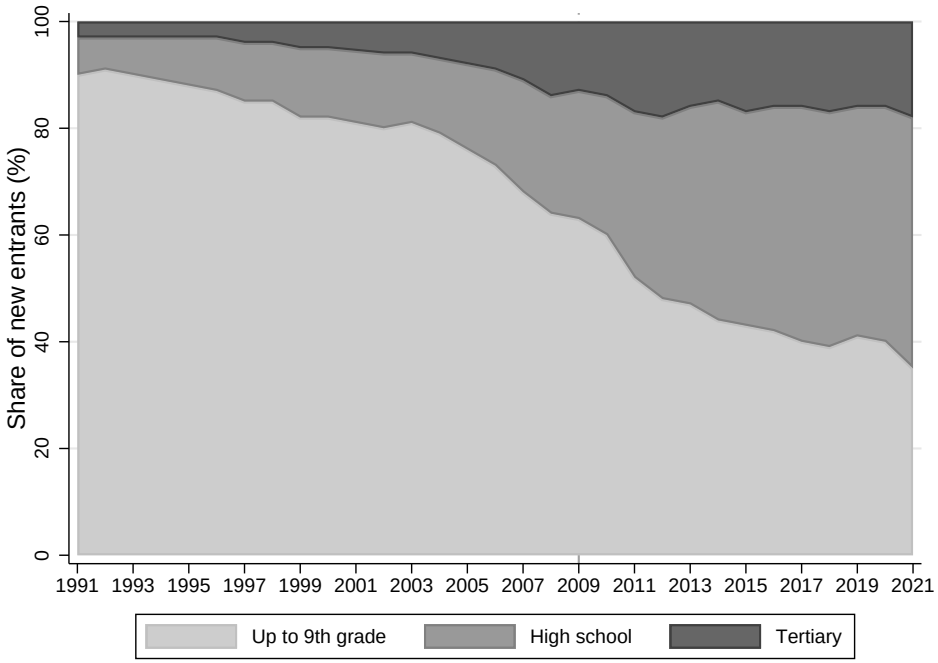
Notes: The dependent variable is the logarithm of the number of students enrolled in upper secondary education (10th–12th grades), considering only programmes aimed at young students and excluding adult-education programmes. The sample is a NUTS 2 region-year panel for mainland Portugal from 2004 to 2023. Each column reports a simple descriptive specification with region fixed effects, region-specific post indicators, and the logarithm of the resident population aged 15–18; heteroskedasticity-robust standard errors are in parentheses. Model 1 codes the post period from 2009 onwards, while Model 2 codes it from 2012 onwards. The near-unit R^2 reflects the mechanical relationship between enrolment counts and the target-age population that generates them, together with the region fixed effects. Because the specification uses only five regions and serves as contextual support, it is presented as descriptive rather than as a stand-alone causal design. *Source:* DGEEC and INE data as documented in Direção-Geral de Estatísticas da Educação e Ciência (DGEEC) (2026) and Instituto Nacional de Estatística (INE) (2026), covering 2004–2023 for mainland NUTS 2 regions.

3.4 Data and descriptive statistics

3.4.1 Data sources

The analysis combines three Portuguese administrative sources linked through a common anonymized firm identifier to build a firm-level panel of manufacturing firms: *Quadros de Pessoal* (*QP*, INE, 2021c), *Comércio Internacional* (*CI*, INE, Statistics Portugal, 2025, and the *Sistema de Contas Integradas das Empresas* (*SCIE*, INE, 2021b). This linked structure makes it possible to observe, within a common panel, both firms' pre-reform exposure to underage labour and their post-reform adjustment along the automation, ICT, and exit margins.

The core source is *QP* (INE, 2021c), a mandatory annual census of all private-sector establishments employing at least one wage earner, administered by the *Gabinete de Estratégia e Planeamento* (GEP)



Notes: The figure plots the annual education shares of new manufacturing entrants. "New entrants" are workers entering the Portuguese private labour market for the first time and observed in manufacturing. Education groups are up to 9th grade, high school, and tertiary. The year 2001 is absent because *Quadros de Pessoal* was not collected that year. *Source:* Authors' computations using *Quadros de Pessoal* (QP) microdata (INE, 2021c).

Figure 3: Education composition of new manufacturing entrants, 1991–2021

of the Ministry of Labour, Solidarity, and Social Security. Because filing is compulsory rather than voluntary, coverage is close to the universe of private-sector employment. In recent waves, the dataset covers roughly 350 thousand firms and 3 million workers per year. The reference month is October from 1994 onward (March before 1994), and data are available from 1986 onward except for 1990 and 2001, when no collection took place. A central advantage of *QP* is that wages are reported by employers rather than self-reported by workers. The data therefore contain a detailed decomposition of earnings – base salary, regular and irregular benefits, overtime pay, and normal hours worked alongside worker characteristics such as age, gender, education, tenure, and occupation, as well as firm-level information on sector, region, firm age, sales, and employment. The anonymized worker and firm identifiers also allow longitudinal tracking over time. In the harmonized dataset, the education and sector variables are taken directly from the underlying *QP* spine. The occupation-based ICT outcome likewise comes from *QP*, but because it requires harmonization across the pre-2010 and post-2010 occupational classifications, we interpret it as a carefully constructed proxy for complementary technical employment rather than as a frictionless measure of digital labour demand. The same occupation records also allow us to construct two occupation-based robot measures, labelled occupation proxy (strict) and occupation proxy (flexible). These firm-year indicators equal one when at least one linked worker is flagged in the corresponding harmonized robot-related occupation mapping. They measure robot-related occupational presence rather than capital imports, but they are part of the chapter's robot-measurement layer alongside the CI-based import measures.

We then use *CI* (INE, Statistics Portugal, 2025), maintained by Statistics Portugal (INE), to measure firm-level robot adoption through import transactions recorded at the 8-digit Combined Nomenclature (Combined Nomenclature (CN)8) product level. *CI* combines two reporting systems: *Intrastat*, which covers intra-EU trade above a reporting threshold, and *Extrastat*, which records extra-EU trade through customs declarations with near-universal coverage. Because the Intrastat threshold changes over time, smaller intra-EU importers may be measured with error in some years. That concern is likely less severe for robot transactions, which are typically high-value capital purchases and may also be sourced from outside the European Union, but it cannot be fully ruled out. Following Acemoglu et al. (2020, 2023), Acemoglu and Restrepo (2022a), and Bonfiglioli et al. (2024), we define industrial robots using the first six digits of the product code, 847950, corresponding to Harmonized System heading 8479.50 ("industrial robots, not elsewhere specified or included"). This classification captures multipurpose, automatically controlled, reprogrammable machines consistent with the ISO 8373 definition, while excluding parts, controllers, and numerically controlled machine tools classified elsewhere. Given Portugal's very limited domestic robot production, imports provide a practical firm-level proxy for robot adoption. At the same time, the measure records the year of import rather than the year the robot enters production, so adoption timing may be measured with some lag. We organize the import-based outcomes in three tiers. Tier 1 is the narrow baseline and uses only Combined Nomenclature six-digit (CN6) code 847950. Tier 2 adds closely related robotic-equipment headings that plausibly accompany robot adoption, while Tier 3 adds Tier 2 plus Computer Numerical Control (CNC) – related automation headings. Taken together with the two occupation-based proxies from *QP*, these form the five robot measures used in the chapter. Table 14 reports descriptive statistics for all five. In the main results, the annual comparison is organised around four outcomes shown side by side - Tier 1 imports, Tier 2 imports, the flexible occupation proxy, and ICT employment - and is reported under the binary, continuous-share, and percentile-share treatment codings. Tier 3, the strict occupation proxy, and the cumulative import measures are retained as supplementary appendix evidence rather than as main-text outcomes.

Finally, we use *SCIE* (INE, 2021b), an enterprise-level dataset produced by INE from the mandatory *Informação Empresarial Simplificada* (IES) filing that firms submit annually to INE, the tax authority, Banco de Portugal, and the commercial registry. *SCIE* is available from 2004 onward and is linked to *QP* and *CI* through the same anonymized firm identifier. We use it primarily to construct the firm-exit indicator from the last year of recorded activity, following the Eurostat/OECD business-demography methodology, and to recover accounting variables used in secondary analysis. Monetary firm-level variables used in the chapter – including sales, robot imports, trade values, value added, labour costs, and investment aggregates – are deflated to constant 2024 euros using the INE consumer-price deflator.

The analysis is restricted to manufacturing firms. In practice, this means that the paper draws on the long time coverage of *QP*, the robot-import information available in *CI* from the mid-1990s onward, and the *SCIE*-based exit and accounting information available from 2004 onward, with the exit outcome discussed from 2006 forward. Longer-run descriptive series use the widest window allowed by each source, while the main estimation samples are restricted to 2002–2019 as described in Section 3.5. Table 12 summarizes

the main variables used in the analysis.

3.4.2 Descriptive statistics

Table 14 is organised around the estimation samples used in the results section rather than around adopters versus non-adopters. Panel A reports firm-year moments for the two main DiD samples: the post-2009 panel contains 435,898 firm-year observations from 35,194 firms, while the post-2012 panel contains 362,557 firm-years from 24,562 firms. The baseline binary exposure equals one in 11.2% of Panel A firm-years and 12.0% of Panel B firm-years. The underlying continuous exposure measure is highly concentrated at zero, which is why the empirical analysis gives priority to the binary exposure, treats the continuous share as the main intensity margin, and also reports percentile bins as a discretized intensity check in the main annual table. The same table also shows how sharply support differs across the chapter's technology-related outcomes. Under Tier 1, the annual robot-importer indicator averages only 0.13–0.15% of firm-years. Tier 2 and Tier 3 are broader but still uncommon, with annual importer means of 0.34–0.39% and 0.63–0.72%, respectively. The occupation-based layer is much less sparse: the flexible proxy is present in 8.4–9.2% of firm-years, while ICT employment is observed in 4.4–4.8% of firm-years. The SCIE-based exit rate is 1.1% in Panel A and 0.6% in Panel B. The relevant descriptive point is therefore not just that the strict robot-import outcome is sparse, but that support differs across the chapter's outcomes and that the occupation-based layer is much less sparse than Tier 1.

Panel B of Table 14 collapses Panel A firms to one observation using their 2006–2008 averages and shows that exposure is not random. Firms that employed minors before the reform are much larger, with average employment of 44.8 workers against 14.3 among less-exposed firms, and they also record higher sales and value added together with a lower labour-cost share. They are more likely to employ ICT workers before the reform (6.5 versus 3.4% of firms), and the flexible occupation proxy shows the same directional selection (15.2 versus 5.9%). On the automation side, all three import tiers remain small in levels before the reform, especially the strict Tier 1 measure, but more-exposed firms do show higher average Tier 2 and Tier 3 importer rates and cumulative import measures. The relevant descriptive message is therefore not that the treated group was already heavily automated, but that exposure loads on scale, organisation, workforce composition, and a broader set of technical occupations. Any empirical interpretation must accordingly come from differential within-firm changes after the reform rather than from raw cross-sectional level comparisons.

The appendix descriptives reinforce the same point and also clarify why Figure A1 and Tables A1–A4 contain many zeros and occasional undefined cells. Those appendix descriptives use the strict Tier 1 baseline, which is extremely sparse even in the full estimation panel: within Panel A, the annual Tier 1 importer margin turns positive in only 572 firm-years across 313 firms. Tier 2 is broader, with 1,500 positive importer firm-years across 674 firms in Panel A and 1,425 across 638 in Panel B, while Tier 3 records 2,725 and 2,609 positive firm-years across 1,071 and 1,009 firms. Appendix Table A1 contrasts firms entering before and after the reform in the full descriptive panel and shows later cohorts to be

Table 12: Variable descriptions

Variable	Description	Source
Strict robot importer (Tier 1)	Indicator equal to 1 if the firm imports CN6 847950 in year t.	CI
Strict cumulative robot imports (Tier 1)	$\log(1 + \text{cumulative real strict robot imports}_{it})$.	CI
Strict annual robot imports (Tier 1)	Annual imports of CN6 847950, deflated to 2024 euros.	CI
Broad robot importer (Tier 2)	Indicator equal to 1 if the firm imports the Tier 1 code set plus closely related robotic-equipment headings.	CI
Broad cumulative robot imports (Tier 2)	$\log(1 + \text{cumulative real broad robotic-equipment imports}_{it})$.	CI
Broad annual robot imports (Tier 2)	Annual imports of the Tier 2 robotic-equipment code set, deflated to 2024 euros.	CI
Automation importer (Tier 3)	Indicator equal to 1 if the firm imports the Tier 2 code set plus CNC-related automation headings.	CI
Automation cumulative imports (Tier 3)	$\log(1 + \text{cumulative real automation-related imports}_{it})$.	CI
Automation annual imports (Tier 3)	Annual imports of the Tier 3 automation and CNC code set, deflated to 2024 euros.	CI
ICT worker	Indicator equal to 1 if the firm employs at least one worker in the harmonized ICT occupation measure.	QP
Occupation proxy (strict)	Indicator equal to 1 if at least one linked worker is flagged in the strict occupation-based robot proxy reconstructed from QP occupation codes.	QP
Occupation proxy (flexible)	Indicator equal to 1 if at least one linked worker is flagged in the flexible occupation-based robot proxy reconstructed from QP occupation codes.	QP

Table 13: Variable descriptions (continued)

Variable	Description	Source
Minor share	Share of workers under age 18 in the firm's QP-linked workforce in year t.	QP
Minor exposure (binary)	Indicator equal to 1 if the firm employed at least one worker under age 18 during 2006–2008.	QP
Minor exposure (share)	Average share of workers under age 18 during 2006–2008.	QP
Minor exposure (percentile category)	Five-category grouping of the 2006–2008 minor-share distribution, with most mass at zero.	QP
Employment	Number of QP worker links observed for the firm in year t.	QP
Average education	Average years of schooling among workers linked to the firm in year t.	QP
Hourly wage	Average hourly wage among linked workers, measured in 2024 euros.	QP
Sales	Firm sales in 2024 euros.	QP / SCIE
Value added	Firm gross value added in 2024 euros.	SCIE
Labour cost share	Personnel expenses divided by value added when value added is positive.	SCIE
Exit	Indicator based on the last observed active year in SCIE following the business-demography rule used in the pipeline.	SCIE

Notes: QP stands for Personnel Records, CI for International Trade records, and SCIE for the Integrated Business Accounts System. The strict industrial-robot measure uses CN6 847950. Tier 2 adds CN6 847989, 851521, 851531, 851539, and 851580. Tier 3 adds Tier 2 plus CNC-related CN6 845610, 845620, 845630, 845710, 845720, 845730, 845811, 845891, 845921, 845931, 845961, 846011, 846021, 846031, and 846241. ICT worker equals 1 if, in firm-year t, the firm employs at least one worker in harmonized ICT occupations (ISCO-08 major groups 25 and 35, strict definition), and 0 otherwise. Occupation proxy (strict) equals 1 if at least one linked worker has CPP2010 raw code 3139 (missing before 2010); occupation proxy (flexible) equals 1 if at least one linked worker has CNP94 raw code 2143, 3114, 3119, 3123, or 8172, or CPP2010 raw code 2151, 3114, 3115, 3119, or 3139. Minor exposure (percentile category) is a five-category measure based on percentile cutoffs of the 2006-2008 minor-share distribution, with firms at zero minor share assigned to the lowest category. Derived interaction terms and event-time dummies are omitted for compactness. The source column reports the primary administrative source for each variable.

smaller, more educated, and less reliant on minors. Tables A2–A4 then show, for the strict baseline measure, that robot importing is concentrated among large firms and a small set of sectors. These descriptives document sample composition, treatment skewness, non-random pre-treatment selection, and the varying support of the three automation tiers. They serve as setup for the empirical strategy, not as informal evidence on treatment effects.

3.5 Research design and empirical strategy

The unit of analysis is the firm-year. The empirical shock is Portugal's 2009 compulsory-schooling reform, which applied nationwide and reduced firms' access to underage labour. Because treatment timing is common, the design identifies coefficients by comparing post-reform outcome changes between firms that were more versus less exposed to underage labour before the reform, rather than by exploiting staggered adoption or regional policy rollout. The estimand is therefore a relative reduced-form effect, not an effect against a fully untreated control group. Its interpretation rests on a conditional parallel-trends assumption, and the discussion below keeps that identifying content separate from the broader mechanism narrative.

3.5.1 Identification strategy

The 2009 reform was a nationwide policy shock, but its bite should have differed across firms depending on their baseline reliance on underage labour. The design therefore exploits cross-firm differences in pre-reform exposure, not differences in the timing or location of the reform itself. In the baseline design, firms that employed at least one worker under age 18 during 2006–2008 are coded as treated, while firms that did not are coded as non-treated. The coefficient of interest is thus the differential post-reform change in outcomes among more exposed firms relative to less exposed firms.

The key identifying assumption is conditional parallel trends: absent the reform, more- and less-exposed firms would have followed comparable outcome paths once time-invariant firm heterogeneity, sector × year shocks, region × year shocks, and baseline education differences interacted with year effects are absorbed. This is a demanding assumption in the present setting. Exposure is predetermined, but it is not random: more exposed firms differ from less exposed firms on size, organisation, and pre-reform workforce composition. The design should therefore be read as a common-timing DiD with non-trivial comparability risk, not as a setting in which firm fixed effects mechanically settle identification.

To capture timing, we exploit the fact that the reform was enacted in 2009 but should only have become fully binding for the first affected cohorts around 2012.[2] We therefore use two post codings. Post2009_t captures enactment and possible anticipation. Post2012_t captures the later phase in which the labour-supply restriction should be more directly visible in firms' hiring environment. These are two

[2]Under the old regime, students could leave school at age 15 after completing 9 years of compulsory education. The 2009 law required them to remain until age 18. The first cohort fully bound by the new rule would have reached labour-market age only a few years later.

Table 14: Descriptive statistics aligned with the empirical design

	Post-2009 sample	Post-2012 sample	
A. Firm-year means in the estimation samples			
Firm-years	435,898	362,557	
Firms	35,194	24,562	
More-exposed share	0.1120	0.1204	
Tier 1 importer	0.0013	0.0015	
Tier 1 annual imports (ths. eur.)	0.3945	0.4511	
Tier 1 cumulative imports (log)	0.0732	0.0863	
Tier 2 importer	0.0034	0.0039	
Tier 2 annual imports (ths. eur.)	0.9620	1.1248	
Tier 2 cumulative imports (log)	0.1749	0.2029	
Tier 3 importer	0.0063	0.0072	
Tier 3 annual imports (ths. eur.)	2.3590	2.7724	
Tier 3 cumulative imports (log)	0.3349	0.3856	
ICT worker	0.0435	0.0476	
Occupation proxy (strict)	0.0451	0.0461	
Occupation proxy (flexible)	0.0838	0.0917	
Minor share	0.0031	0.0030	
Exit (SCIE sample)	0.0112	0.0064	
B. Pre-reform firm means by binary exposure (2006–2008, Panel A firms)			
	Non-treated	Treated	Difference
Firms	31,734	3,460	
Minor share	0.0000	0.0534	0.0534
ICT worker	0.0338	0.0653	0.0315
Occupation proxy (strict)	.	.	.
Occupation proxy (flexible)	0.0587	0.1520	0.0933
Tier 1 importer	0.0007	0.0006	-0.0002
Tier 1 annual imports (ths. eur.)	0.4048	0.1037	-0.3011
Tier 1 cumulative imports (log)	0.0293	0.0915	0.0622
Tier 2 importer	0.0023	0.0040	0.0017
Tier 2 annual imports (ths. eur.)	0.9706	0.3734	-0.5972
Tier 2 cumulative imports (log)	0.0862	0.2281	0.1420
Tier 3 importer	0.0041	0.0095	0.0055
Tier 3 annual imports (ths. eur.)	1.6841	3.5389	1.8548
Tier 3 cumulative imports (log)	0.1732	0.4591	0.2859
Average education (years)	7.223	7.061	-0.162
Hourly wage (2024 euros)	4.838	4.824	-0.014
Employment	14.317	44.836	30.519
Labour cost share	1.319	0.848	-0.472
Sales (thous. 2024 euros)	2,406.6	3,858.5	1,451.9
Value added (thous. 2024 euros)	755.7	1,324.0	568.3

Notes: Panel A reports firm-year means for the two main estimation samples over 2002–2019. The robot-measurement layer consists of five variables: Tier 1 is the strict CN6 847950 robot measure; Tier 2 adds closely related robotic-equipment headings; Tier 3 adds CNC-related automation headings; and "Occupation proxy (strict)" and "Occupation proxy (flexible)" are QP-based indicators that equal one when the firm-year contains at least one worker flagged in the corresponding reconstructed occupation-based robot proxy. The "Exit" row uses the SCIE-linked exit samples, available from 2006 onward. Panel B collapses the post-2009 estimation sample to one observation per firm using averages over 2006–2008 and is intended to document pre-treatment selection rather than post-treatment differences. "Treated" denotes firms that employed at least one worker under age 18 during 2006–2008; "Difference" is treated minus non-treated. The table is descriptive only and should not be read as a formal balance test.

Source: Authors' calculations using QP, CI, and SCIE data, 2002–2019 (Panel A) and 2006–2008 (Panel B).

timing windows around the same national reform, not separate treatments or separate quasi-experiments. Stronger effects under the later coding are consistent with delayed firm adjustment, but any such strengthening may appear through internal workforce reorganization or complementary technical employment rather than through robot imports alone.

Mechanism and scope. The economic mechanism of interest runs from the reform to lower underage employment in more exposed firms, from lower underage employment to tighter access to a specific low-cost labour input, and from that tightening to firm adjustment. The chapter documents the first link directly through firm-level first-stage regressions and studies the latter links in reduced form across three margins: narrow robot imports, occupation-based workforce reorganization, and complementary ICT employment. The question is whether firms that were more exposed to the loss of underage labour adjusted differently after the reform, not whether this chapter fully closes the broader mechanism or identification debate.

3.5.2 Firm-level specifications

For orientation only, the appendix reports a pooled DiD benchmark without firm fixed effects. The main specification is a firm-level fixed-effects model with flexible sectoral and regional time shocks:

$$\begin{aligned} Y_{it} = \beta(\text{Post}_t \times \text{Exposed}_i) \\ + \sum_{\tau} \theta_\tau \left(\text{Educ2008}_i \times \mathbf{1}[t = \tau]\right) \\ + \sum_{\tau} \rho_\tau \left(\text{Size2008}_i \times \mathbf{1}[t = \tau]\right) \\ + \gamma_i + \lambda_{s(i)t} + \mu_{r(i)t} + \epsilon_{it} \end{aligned} \tag{5}$$

Here, Y_{it} is the outcome for firm i in year t; Exposed_i is a time-invariant pre-reform exposure measure; Post_t is either Post2009_t or Post2012_t; Educ2008_i is the firm's baseline average years of education; Size2008_i is the firm's baseline log employment; γ_i are firm fixed effects; $\lambda_{s(i)t}$ are sector × year effects; and $\mu_{r(i)t}$ are NUTS 2 region × year effects. The main coefficient of interest is β. Standalone exposure and post terms are absorbed by the fixed effects. Standard errors are clustered at the firm level, which is the level at which exposure is defined. Under this specification, β captures a within-firm post-reform outcome change that is larger in more exposed firms than in less exposed firms after netting out common sectoral and regional time shocks. For brevity, the chapter refers to this as the main firm fixed-effects TWFE benchmark, even though the time effects are allowed to vary flexibly by sector and region rather than entering only as common year fixed effects.

Because treatment timing is common rather than staggered, the usual negative-weighting concern in staggered-adoption TWFE designs is not the first-order issue here (Callaway & Sant'Anna, 2021; de Chaisemartin & D'Haultfœuille, 2020). The central issue is instead whether more- and less-exposed firms were already on different trajectories. For that reason, the later diagnostics emphasize timing,

trend differences, and comparability rather than treating a modern staggered estimator as an alternative identification strategy.

Event-study specification. To assess timing and pre-trends, we replace the single post indicator with year-by-year interactions between the binary exposure indicator and calendar-year dummies:

$$\begin{aligned} Y_{it} = & \sum_{k \neq 2008} \beta_k \left(\text{Exposed}_i \times \mathbf{1}[t = k]\right) \\ & + \sum_{\tau} \theta_\tau \left(\text{Educ2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \sum_{\tau} \rho_\tau \left(\text{Size2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \gamma_i + \lambda_{s(i)t} + \mu_{r(i)t} + \epsilon_{it} \end{aligned} \tag{6}$$

The reference year is 2008. Under conditional parallel trends, the pre-reform interaction coefficients should be jointly close to zero. The main text reports this diagnostic for the baseline binary exposure in Panel A, which provides the longer pre-reform window, and the appendix extends the same exercise to the broader outcome definitions. We do not re-cast the event study around alternative share thresholds or transformed exposure-intensity measures: those codings are useful for descriptive splits and static sensitivity checks, but in a dynamic specification they would either impose arbitrary cutoffs or turn the plotted coefficient into a slope with respect to an exposure distribution that is heavily concentrated at zero. The event study is informative because it speaks directly to timing in the maintained design; it is not presented as a blanket validation device.

First-stage specification. To test whether the reform reduced underage employment in more exposed firms, we estimate:

$$\begin{aligned} \text{MinorOutcome}_{it} = & \beta(\text{Post}_t \times \text{Exposed}_i) \\ & + \sum_{\tau} \theta_\tau \left(\text{Educ2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \sum_{\tau} \rho_\tau \left(\text{Size2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \gamma_i + \lambda_{s(i)t} + \mu_{r(i)t} + \epsilon_{it} \end{aligned} \tag{7}$$

where MinorOutcome_{it} is alternately the share of workers under 18 in firm i's workforce in year t, the number of minors M_{it}, or an indicator for employing any minor. These first-stage regressions use the same exposure definitions, post codings, sample splits, fixed effects, controls, and clustering structure as the main firm-level design, changing only the dependent variable. Their role is necessary but limited: they show whether the reform reduced access to underage labour in firms that had relied on it before the reform, but they do not by themselves resolve the broader comparability or timing concerns.

3.5.3 Outcomes, exposure, and sample construction

Outcomes. The outcome set is organized around three adjustment margins. The first is import-based robot adoption, captured by annual Tier 1 and Tier 2 importer indicators together with cumulative import measures, where Tier 1 is the strict industrial-robot import heading CN6 847950 and Tier 2 adds closely related robotic-equipment headings. Tier 3 broadens the import concept further to automation and CNC)-equipment headings. The annual importer outcomes ask whether more exposed firms are more likely to record a robot-related import in a given year. Because robot purchases are lumpy and the year of import need not coincide exactly with the year in which the equipment is integrated into production, we also estimate cumulative import outcomes that sum observed robot-related imports within firm over time and transform them as $\log(1 + \text{cumulative imports})$ in the baseline, with an inverse-hyperbolic-sine robustness check reported in B.2. These cumulative outcomes ask a different question from the annual importer margins: not whether exposed firms buy in a particular year, but whether they end up with a larger accumulated stock of robot-related imports over the post-reform period. The second margin is internal workforce reorganization, captured by two occupation-based firm-year indicators from QP - occupation proxy (strict) and occupation proxy (flexible) - built from harmonized robot-related occupation mappings. The flexible proxy receives more emphasis in the results because it is available on the full annual panel, is materially less sparse than the import measures, and has cleaner timing diagnostics than the strict annual importer in the current event-study exercise. The strict proxy is retained as a supplementary comparison rather than as a main-text outcome. The third margin is complementary technical employment, captured by an ICT-employment indicator constructed from harmonized occupation codes. ICT is not itself a robot measure; it captures a complementary technical-labour margin that is observed much more frequently than strict robot importing. The annual comparison in Table 16 focuses on Tier 1, Tier 2, the flexible occupation proxy, and the presence of ICT workers. Table 16 reports these four outcomes under the binary exposure, the continuous share, and the percentile-share coding to show whether the ranking depends on how exposure is parameterized. The cumulative margins remain in Appendix B.2, where they are treated as stock-based sensitivity checks rather than as main estimates. A separate secondary outcome is firm exit from the SCIE-linked sample. Because exit is an absorbing binary outcome observed only from 2006 onward and on a narrower support than the workforce and trade outcomes, we discuss it later in a dedicated subsection and give it less evidentiary weight than the four main annual adjustment outcomes. For consistency with the rest of the chapter, the exit subsection still reports the binary, continuous-share, and percentile-share exposure codings, but its main specification remains the firm fixed-effects linear probability model.

These measures should not be treated as interchangeable objects. Tier 1–Tier 3 import outcomes capture equipment acquisition in international markets and are the sparsest part of the design. The annual Tier 1 importer is conceptually narrow but empirically thin. The cumulative import outcomes help address the concern that annual purchases are too lumpy to reveal a gradual stock response, but they are not cleaner causal margins: because they accumulate past imports within firm, they inherit support and composition

issues from earlier periods and also carry the usual interpretation problems of the $\log(1 + Y)$ transformation when the outcome is frequently zero (Chen & Roth, 2024). For that reason, the cumulative specifications serve as informative stock-based sensitivity checks rather than as main outcomes. The occupation-based proxies capture something different: internal occupational presence associated with automation-related tasks rather than imports of robotic equipment. ICT likewise captures complementary technical labour rather than installed robot capital. The import, occupation, and ICT margins therefore capture different forms of firm adjustment rather than interchangeable automation measures.[3]

Exposure. Exposed_i is constructed from firms' employment of minors during the three pre-reform years 2006–2008. The baseline exposure is binary: equal to one if the firm employed at least one worker under 18 in that period, and zero otherwise. This is the baseline treatment measure because it imposes the fewest functional-form assumptions and yields the clearest more-exposed-versus-less-exposed contrast. We also report two alternative codings: the continuous pre-reform share of minors in the firm's workforce, and a percentile-based categorical version that leaves the mass at zero in the bottom category and then partitions the upper tail of the exposure distribution into progressively more exposed bins. The continuous share is the chapter's main intensity specification. The percentile-based coding is more discretized and is reported in the main annual comparison as a transparent functional-form check, but it carries less interpretive weight than the binary and continuous-share codings. The main descriptive figure at the start of Section 6 uses the same baseline binary exposure for comparability with the maintained design, while the stricter threshold-based "more exposed - treated" versus "less exposed - non-treated" splits shown in the appendix descriptive figures are only visual devices for raw comparisons and are not part of the maintained regression design.

Timing and sample construction. We estimate the design under two post codings:

- $\text{Post2009}_t = 1$ from 2009 onward, capturing enactment and possible anticipation.
- $\text{Post2012}_t = 1$ from 2012 onward, capturing the phase in which the labour-supply restriction should have become most directly binding.

These two post codings are alternative timing windows around the same national reform, not separate treatments. The main estimation window is 2002–2019. We begin in 2002 to place some distance between the current reform and the earlier 1993 compulsory-schooling extension while retaining a meaningful pre-period, and we end in 2019 to avoid contamination from the COVID-19 shock. The analysis is restricted to manufacturing firms. Outcome-specific samples can begin later when the required CI or SCIE information is only available from the mid-2000s onward. In practice, the SCIE-linked exit models begin

[3] The occupation mapping bridges the pre-2010 and post-2010 coding systems and is designed to isolate ICT-related occupations as consistently as the available classifications allow.

in 2006, while the main workforce- and import-based panels are defined over the common 2002–2019 window.

We construct two panels. Panel A conditions on firms that existed in 2009 and Panel B conditions on firms that existed in 2012. In the manufacturing analysis panel, these samples contain 435,898 firm-year observations from 35,194 firms (Table 14); the main four-outcome TWFE table uses 428,076 and 353,203 firms firm-year observations after drops due to missing baseline controls and singletons. Because exact counts vary slightly across outcomes and estimators, the tables report regression-specific observations. Panel B is better aligned with the later binding phase of the reform, but it mechanically excludes firms that disappeared between 2009 and 2012 and therefore embeds an additional selection margin. The two panels should therefore be read as sample definitions aligned with the two post codings, not as separate identification strategies.

3.5.4 Heterogeneity and mechanisms

The heterogeneity analysis is used to locate where any average differential response comes from, not to define a separate identification strategy. Because the binary exposure indicator is correlated with firm size, these exercises are interpretive extensions around the same baseline design rather than a separate source of identification. For consistency with Table 16, we keep the same four annual outcomes - Tier 1 importing, Tier 2 importing, the flexible occupation proxy, and ICT employment - but we simplify the treatment coding and report heterogeneity only for the baseline binary exposure. Once the sample is split into size and labour-intensity subsamples, the continuous-share and percentile treatments add noise faster than they add identification. We therefore use the subgroup results to ask where the main annual responses are concentrated, not to create a second layer of treatment variation. We re-estimate the baseline TWFE specification by quartile of the pre-reform firm-size distribution and by pre-reform labour intensity, measured using the baseline labour-cost share from SCIE. These exercises are especially useful in the present chapter because the more stable positive margins in Section 3.6.3 appear concentrated in firms with greater organizational capacity.

To formalise the size gradient, we also estimate a triple-difference specification that interacts the baseline DiD term with a large-firm indicator:

$$\begin{aligned} Y_{it} = {} & \beta_1(\text{Post}_t \times \text{Exposed}_i) + \beta_2(\text{Post}_t \times \text{Exposed}_i \times \text{Large}_i) \\ & + \sum_{\tau} \theta_\tau \left(\text{Educ2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \sum_{\tau} \rho_\tau \left(\text{Size2008}_i \times \mathbf{1}[t = \tau]\right) \\ & + \gamma_i + \lambda_{s(i)t} + \mu_{r(i)t} + \epsilon_{it} \end{aligned} \tag{8}$$

Here, Large_i is an indicator for above-median pre-reform employment. The coefficient β_1 captures the DiD effect for smaller exposed firms, while β_2 is the incremental effect for larger exposed firms. This

specification is informative about where the average effect is concentrated, but it does not resolve the underlying comparability problem because firm size is itself endogenous to broader firm characteristics.

3.5.5 Supporting diagnostics and supplementary estimators

The appendix reports robustness checks and diagnostic tests relevant for interpreting the results. The most informative are the event study, the age-based and fake-timing placebos, the trend-augmented TWFE specification, and the matched common-support design. These exercises do not replace the baseline specification, but they show where the evidence is stable and where it becomes more fragile.

The event study is especially useful because it speaks directly to timing in the maintained design. The placebo exercises then ask whether similar patterns appear where the substantive channel should be absent, using both negative-control exposure measures and fake reform years. That matters in this setting because treatment timing is common but exposure is not random.

Trend-augmented TWFE and the matched common-support design ask a different question: whether the baseline coefficients survive when the comparison between more- and less-exposed firms is made more demanding. Trend-augmented TWFE adds an exposure-specific linear time trend, $\text{Exposed}_i \times (t - 2008)$, allowing the two groups to follow different linear trajectories even absent the reform. The matched design first restricts the sample to exact common-support cells defined by baseline sector, region, and firm-size class, and then reweights less-exposed firms within those cells using pre-2009 outcomes and firm characteristics. These exercises do not solve selection on unobservables, but they tighten the comparison set in a way that is directly relevant for the chapter's main concern.

Because treatment timing is common, we also report a single-cohort imputation Average Treatment on the Treated (ATT) for the binary exposure. More exposed firms are coded as first treated in 2009 in Panel A and 2012 in Panel B, while less exposed firms are coded as never treated. This is a common-timing estimator check rather than an alternative identification strategy.

Finally, we report a narrow PPML sensitivity for the binary robot outcomes in levels and a comparison of firm-level and sector–region cell-level clustered standard errors. The reason for using PPML here is functional-form rather than identification: for non-negative and highly skewed outcomes, it avoids the $\log(1 + Y)$ transformation and estimates the conditional mean directly in levels. In the current panel, however, this check is materially narrower than the main design. To keep estimation feasible with firm fixed effects on extremely sparse outcomes, the PPML regressions absorb firm and common year effects rather than the full sector $\times$ year and region $\times$ year structure of Equation 5, and `ppmlhdfe` drops firms with separated all-zero outcome paths, leaving a much smaller retained sample. The PPML coefficient is therefore interpreted as a multiplicative effect on the conditional mean in that retained sample – for the importer outcome, on the conditional probability of importing; for cumulative robot value, on expected cumulative import value in levels. We use this diagnostic only to ask whether a level-based non-linear mean specification overturns the sign ranking of the linear baseline. It is not a substitute for the chapter's main DiD coefficient. The clustering comparison is likewise auxiliary: it asks whether the main inference

is sensitive to a coarser dependence structure, but it does not address the chapter's main design risks, namely timing, differential trends, comparability, spillovers, and measurement.

3.5.6 Attrition and aggregate checks

Two secondary exercises address specific concerns that the firm-level design leaves open. First, because the main samples are unbalanced and more exposed firms may exit at different rates, we re-estimate the main specification on balanced-panel restrictions (firms observed for at least 12 years and firms observed for all 18 panel years). This tests directly whether differential attrition drives the main results. Second, in the appendix we aggregate firms to manufacturing sector × NUTS 2 region cells and estimate a coarser DiD with cell and year fixed effects:

$$Y_{jt} = \alpha + \beta_1(\text{Post}_t \times \text{Exposed}_j) + \beta_2(X_j \times \text{Year}_t) + \gamma_j + \lambda_t + \epsilon_{jt} \tag{9}$$

where Y_{jt} is the outcome for cell j in year t and cell exposure is a dummy for high minor intensity (pre-reform cell share above 14%). The cell-level outcomes mirror the four firm-level margins. This specification is much weaker than the firm-level design – it is estimated on far fewer units and does not condition on within-cell firm heterogeneity – so it is reported as an appendix complement rather than as a main-text exercise.

3.5.7 Threats to identification

Comparability. The first-order issue is whether more- and less-exposed firms are comparable enough for the DiD contrast to be informative. Pre-reform exposure to underage labour is not randomly assigned. Firms classified as more exposed under the 2006–2008 binary measure are larger and organisationally different from less-exposed firms even before the reform, so firm fixed effects and flexible controls do not guarantee that their counterfactual paths would otherwise have remained comparable. The coefficients should therefore be read as relative reduced-form effects of greater pre-reform exposure, not as structural elasticities or as estimates against a fully unaffected control group.

Crisis overlap and spillovers. The reform also overlaps with the sovereign debt crisis and with broader restructuring in Portuguese manufacturing. If firms that relied more on minors were also more exposed to crisis-era contraction, part of the estimated response may reflect crisis-induced adjustment rather than the reform itself. Because the reform is national, less-exposed firms are not a pure untreated group either: general-equilibrium or local-labour-market spillovers may affect them as well. These concerns do not invalidate the design, but they narrow what can be claimed from it.

How far the ICT result can be taken. The main positive downstream result is on ICT employment, and that is also where the diagnostic evidence bites hardest. In the baseline binary design, the ICT

coefficient is 0.0116 in Panel A and 0.0163 in Panel B. But the age-based exposure placebos – which replace minors with young workers aged 19–21 or 20–22 who were not directly affected by the reform – produce ICT coefficients around 0.0234–0.0244, comparable to or larger than the main estimate. The fake-timing exercises also produce small positive pseudo-effects for ICT in some pre-reform years. Designs that tighten comparability then attenuate the result sharply. Allowing more-exposed firms to follow a differential linear trend moves the ICT coefficient to -0.0031 in Panel A and 0.0086 in Panel B. In the matched common-support exercise it falls to 0.0023 and 0.0056, and with added trends to -0.0029 and 0.0085. Taken together, these diagnostics do not erase the ICT pattern, but they make it difficult to interpret as cleanly reform-specific. The import measures face a different problem: they come from trade data, they are extremely sparse on the annual extensive margin, and the cumulative outcomes remain stock-like even under alternative transformations. The occupation-based proxies and the ICT measure are richer for internal adjustment, but they should still be interpreted as distinct outcomes with different strengths and vulnerabilities rather than as interchangeable automation measures.

3.6 Results

We present the results in order of evidentiary strength. The first stage establishes whether the reform tightened the underage-labour margin in more exposed firms. The event study then summarizes the main timing evidence. Section 3.6.3 reports the central firm-level fixed-effects estimates across three margins: import-based robot measures, occupation-based reorganization, and complementary ICT employment. Section 3.6.4 examines heterogeneity by firm size and labour intensity, and Section 3.6.5 studies exit and balanced-panel restrictions. Appendix sensitivities are used to bound interpretation rather than to generate alternative main estimates.

Figure 4 shows the full annual comparison set used in the empirical analysis: the Tier 1, Tier 2, and Tier 3 import indicators, the occupation proxies (strict and flexible), and ICT employment. In this figure, "more exposed" is defined using the same baseline binary exposure as in the maintained regressions, namely whether the firm employed any worker under age 18 during 2006–2008. The strict Tier 1 importer remains too sparse to support much visual interpretation beyond the fact that the annual extensive margin is extremely thin. On the broader annual importer margins, the occupation proxies, and ICT, the more-exposed group is generally above the less-exposed group in raw levels for much of the sample. That pattern matters because it shows that the baseline binary exposure loads not only on reliance on minors but also on broader pre-existing differences in firm size, organization, and technical intensity.

Appendix Figure B2 reports the corresponding cumulative import outcomes under the same binary exposure, and the cumulative patterns reinforce the same point: the more-exposed group generally starts above the less-exposed group and remains above it, with the raw gap widening as the import definition broadens from Tier 1 to Tier 3. Appendix Figures B3 and B4 then show that these raw treated–control comparisons are highly sensitive to how exposure is defined. Relative to the baseline binary split, tighter

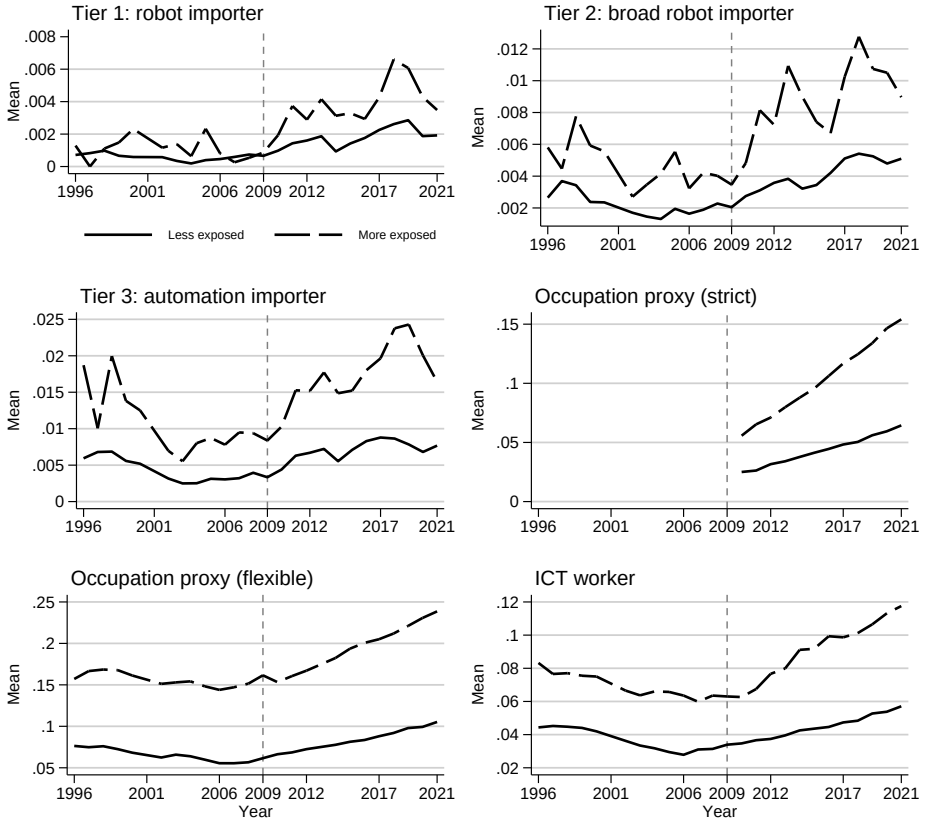


Figure 4: Mean annual robot, automation, occupation-proxy, and ICT outcomes for more exposed and less exposed firms, 1996–2021 (Tier 1, Tier 2, Tier 3, occupation proxies, and ICT; exposure defined as any worker under age 18 during 2006–2008)

share-based cutoffs compress and often reverse the raw ordering on the broader automation, occupation, and ICT margins, while the treated group becomes much smaller. That pattern suggests that the baseline binary definition loads partly on firm size, organisation, and pre-existing technical sophistication, whereas stricter share cutoffs isolate a smaller set of firms that appear more genuinely reliant on underage labour. Two caveats therefore apply. First, the regression panel is restricted to 2002–2019 (see Section 3.5), whereas these descriptive figures extend through 2021. Second, because the level ranking itself changes across exposure definitions, these figures mainly show support, composition, and raw trend differences rather than informal treatment effects. The regression estimates that follow absorb firm, sector–year, and region–year fixed effects and therefore provide the more relevant assessment of differential post-reform adjustment.

3.6.1 First stage: the reform reduced underage employment

Before interpreting any automation or ICT response, we verify that the reform reduced underage employment in more exposed firms. Table 15 reports the binary-exposure first stage using the same specification, fixed effects, and clustering as the main design. Because 96.6% of firm-year observations have no minors,

Table 15: First-stage firm-level DiD estimates for minor employment

	Share of Minors	Number of Minors	Any Minor
Panel A: Post 2009			
Exposed × Post	-0.0315***	-0.4384***	-0.2855***
	(0.001)	(0.011)	(0.004)
Observations	428,076	428,076	428,076
Panel B: Post 2012			
Exposed × Post	-0.0223***	-0.3758***	-0.2401***
	(0.001)	(0.010)	(0.004)
Observations	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes
Baseline educ./size × year	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 7, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Entries report the coefficient on the binary exposure interacted with the relevant post indicator. Columns report the share of workers under 18, the number of minors M_{it}, and an indicator for employing any minor. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Binary exposure equals 1 if the firm employed at least one worker under 18 during 2006–2008. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP data, 2002–2019.

the table reports three versions of the first-stage margin: the share of minors, the number of minors M_{it}, and an indicator for employing any minor. In Panel A, the reform is associated with a 3.15 percentage-point decline in the share of minors, a reduction of 0.44 minors per firm, and a 28.6 percentage-point decline in the probability of employing any minor. In Panel B, the corresponding magnitudes are 2.23 percentage points, 0.38 minors, and 24.0 percentage points. Panel B coefficients are uniformly smaller, consistent with the later post coding and with survivor selection. These are large relative effects – the share of minors falls by more than half in more exposed firms – but the absolute levels are small: even in the pre-reform period, minors represented only about 3–5% of the workforce in treated firms. The labour-supply channel is therefore real but narrow, and the downstream technology-related results should be read accordingly.

3.6.2 Event-study evidence

Figure 5 reports the event study from Equation 6 for Panel A under the binary exposure, with 2008 as the reference year. We focus on Panel A because it provides the longer pre-reform window. The event study uses the same fixed effects, baseline controls, and clustering as the main specification.

The figure compares the full annual outcome set used later in Section 3.6.3: the Tier 1, Tier 2, and Tier 3 importer indicators, the strict and flexible occupation proxies, and ICT employment. Its main role is to show whether the timing evidence differs across outcomes. The annual Tier 1 importer has the weakest

timing case: its pre-reform coefficients are already positive in 2003–2005, and the joint pre-trend null is rejected ($p = 0.016$). ICT and the flexible occupation proxy do not reject the joint pre-trend null ($p = 0.193$ and $p = 0.560$), and neither do Tier 2 and Tier 3 importers ($p = 0.289$ and $p = 0.603$). This is still noisy evidence rather than a clean dynamic treatment profile, but it places ICT and the flexible occupation proxy on firmer timing ground than the strict annual importer. The strict occupation proxy is less informative in this exercise because its narrower support means some pre-reform interactions are absorbed in the current specification.

Appendix Figure B1 reports the cumulative import event studies across the three import tiers. None of these cumulative margins rejects the pre-trend null: $p = 0.493$ for Tier 1, $p = 0.537$ for Tier 2, and $p = 0.737$ for Tier 3. That helps on the narrow timing diagnostic, but the dynamic paths remain flat and noisy, and the interpretation problem is unchanged: cumulative outcomes are stock-like objects that trend mechanically by construction, so passing the pre-trend test is not enough to turn them into clean adoption measures.

In summary, the strict annual importer has the weakest timing case. Tier 2, Tier 3, the flexible occupation proxy, and ICT are more compatible with the parallel-trends assumption in this exercise. The strict occupation proxy remains only partially diagnosable because of its narrower support.

3.6.3 Firm-level adjustment margins: robot imports, occupational reorganization, and ICT employment

Table 16 reports the main firm-level estimates from Equation 5. Appendix Table B1 provides a pooled DiD benchmark without firm fixed effects for comparison. The pooled estimates are uniformly more expansionary – under the binary exposure, the pooled interaction is positive for annual Tier 1 importing, cumulative robot imports, and ICT in both post codings. This reflects persistent cross-sectional differences between treated and control firms rather than a within-firm reform effect, and it is the main reason for preferring the firm fixed-effects specification.

Table 16 reports the main annual estimates from Equation 5, using firm, sector × year, and region × year fixed effects together with baseline education- and baseline size-by-year controls. It places four annual outcomes side by side: the Tier 1 importer, the Tier 2 importer, the flexible occupation proxy, and ICT employment. The binary exposure is the baseline estimand, the continuous pre-reform share of minors is the main intensity specification, and the percentile-share treatment is a discretized functional-form check. Appendix Table B4 adds the cumulative import measures. Those cumulative outcomes matter because robot purchases are lumpy and the year of import need not coincide with the year of operational use. They therefore answer a separate stock question – whether more exposed firms end up with larger accumulated robot-related imports – rather than duplicating the annual table. Supplementary Table B13 extends the annual comparison to Tier 3 and to the strict occupation proxy, and Supplementary Table B15 records the corresponding Panel A event-study lead-test p-values. Taken together, the evidence separates naturally into three margins.

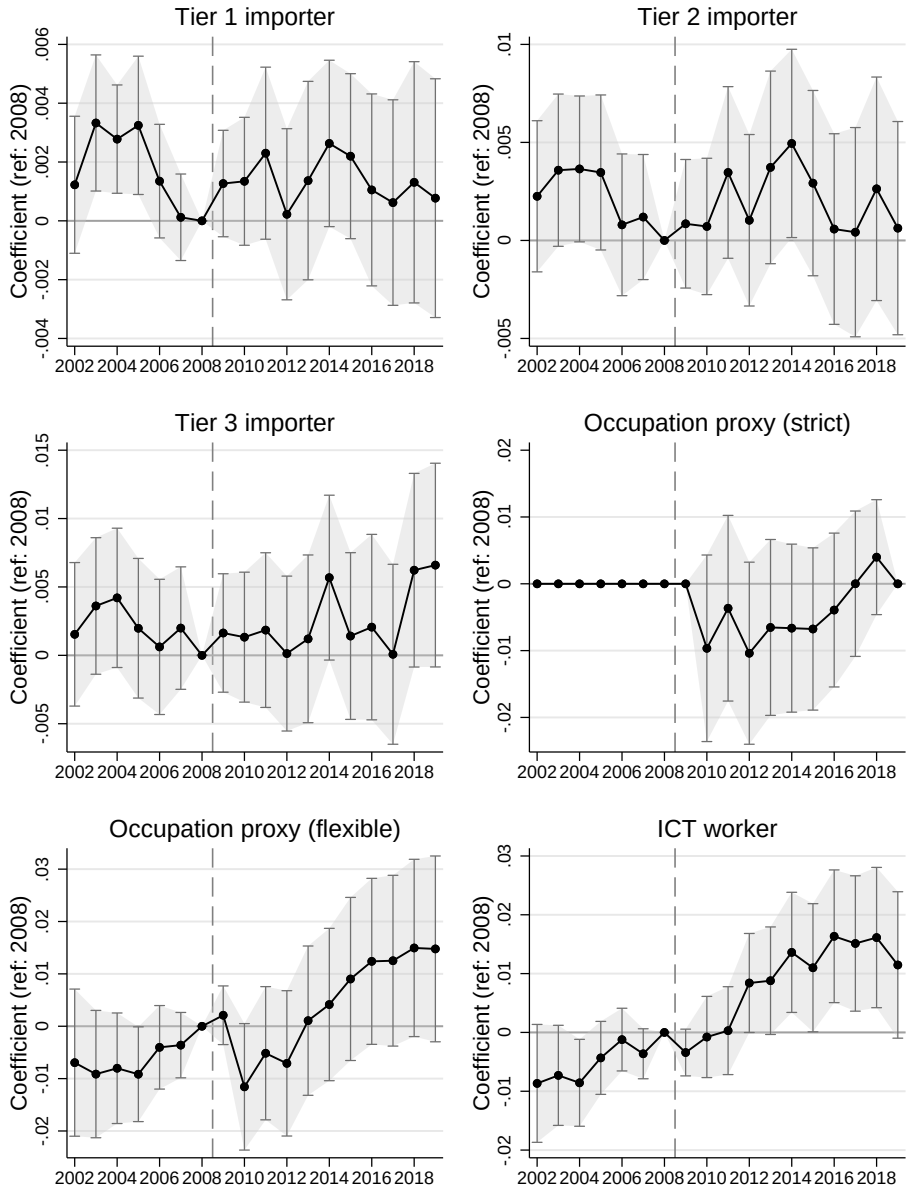


Figure 5: Event study for annual import, occupation, and ICT outcomes under the baseline binary exposure (Panel A)

The comparison between Panel A and Panel B is informative, but not as a clean before-and-after timing exercise. Panel B moves the post window to the later and more binding phase of the reform, so larger coefficients there are compatible with delayed adjustment. But Panel B also conditions on firms surviving through 2012. The useful comparison is therefore whether the ordering of outcomes is stable across the two panels and whether any strengthening appears in the internal workforce margins rather than in the sparse import measures. On that test, the message is stable: the import-based margins remain near zero in both panels, while the flexible occupation proxy and especially ICT are somewhat larger in Panel B.

The first margin is robot imports. On that margin, the annual extensive-margin coefficients are essentially zero throughout the main table. Under the binary exposure, Tier 1 is -0.0002 in Panel A and -0.0003 in Panel B, while Tier 2 is effectively zero in Panel A and only 0.0003 in Panel B. Under the continuous-share exposure, the same coefficients remain small and negative, at -0.0019 and -0.0035 for Tier 1 and -0.0025 and -0.0031 for Tier 2. Under the percentile-share treatment, Tier 1 remains at -0.0001 and -0.0002, while Tier 2 is -0.0002 and -0.0001. Supplementary Table B13 shows that broadening the annual import concept to Tier 3 does not change that conclusion: the Tier 3 coefficients remain small and imprecise at 0.0005 and 0.0009 under the baseline binary design. The cumulative import proxies ask whether exposed firms end up with more accumulated robot-related imports even if annual purchases are too infrequent or mistimed to appear in importer status. They do not uncover a positive pattern either. Cumulative Tier 1 imports are -0.0400 in Panel A and -0.0421 in Panel B under the binary treatment, while cumulative Tier 2 imports are -0.0508 and -0.0466; the continuous-share codings remain negative as well. Appendix Table B5 leaves the same qualitative pattern, with binary coefficients of -0.0422 and -0.0445 for Tier 1 and -0.0535 and -0.0488 for Tier 2, and Supplementary Table B14 shows that widening the cumulative import definition to Tier 3 does not change the conclusion. These cumulative regressions are useful because they address the concern that annual imports may be too lumpy to line up with the reform. But they remain secondary because accumulation makes them stock-like and path-dependent. The conclusion is therefore limited but clear: neither the annual flow margin nor the accumulated stock margin provides robust evidence that more exposed firms increased robot imports after the reform.

The second margin is internal occupational reorganization. Table 16 shows that the flexible occupation proxy is positive under all three main-text treatment codings: it is 0.0082 in Panel A and 0.0127 in Panel B under the binary exposure, 0.0490 and 0.0535 under the continuous-share exposure, and 0.0041 and 0.0051 per percentile bin. This outcome does not measure equipment purchases. It captures whether the firm employs workers in automation-related occupations, so it is closer to internal workforce recomposition than to robot adoption. The flexible proxy therefore provides suggestive, but not decisive, evidence that more exposed firms reorganized their workforce in this direction even though they did not import more robots on a robust basis. Its somewhat larger Panel B coefficients are consistent with gradual internal adjustment, but Panel B also conditions on firms surviving through 2012, so we do not treat that strengthening as a clean delayed-treatment effect. Its Panel A lead-test p-value is 0.560, compared with 0.016 for Tier 1, 0.289 for Tier 2, and 0.193 for ICT (Supplementary Table B15), which places it on firmer timing

ground than the strict import margin without turning it into a clean causal result. We place less weight on the strict occupation proxy, which is available on a narrower support and is less informative in the current panel.

The third margin is complementary ICT employment. ICT is the clearest positive response in the main annual table: under the binary exposure, the coefficients are 0.0116 in Panel A and 0.0163 in Panel B, both statistically significant; under the continuous-share exposure they are 0.0337 and 0.0435; and under the percentile-share treatment they remain positive and precisely estimated at 0.0037 and 0.0047 per exposure bin. This is the chapter's main downstream result. But ICT employment is not a direct measure of automation or robot use. It captures complementary technical hiring, so the result is best read as evidence of differential workforce adjustment. Relative to the import outcomes, the ICT pattern is larger and more stable across the two panels. The somewhat stronger Panel B coefficients are consistent with gradual internal adjustment, but Panel B also conditions on survival through 2012, so that difference is not a clean timing comparison. Nor is the ICT pattern reform-specific in a clean sense: age-based placebo exposures generate ICT coefficients around 0.0234–0.0244, larger than the baseline estimates, while the coefficient falls to -0.0031 and 0.0086 under trend-augmented TWFE, to 0.0023 and 0.0056 in matched TWFE, and to -0.0029 and 0.0085 once matched trends are added. We therefore treat the ICT result as the main downstream pattern, but as bounded evidence on differential workforce adjustment rather than as a clean reform-specific effect.

Appendix Tables B6 and B7 report an alternative control specification that drops baseline size $\times$ year from the control set while retaining baseline education $\times$ year. Under that specification, the binary robot-import coefficients turn positive (Tier 1 rises to 0.0013 in Panel A and 0.0015 in Panel B; Tier 2 to 0.0020 and 0.0026) and the cumulative measures become large and positive. ICT barely changes. The pattern is instructive: because pre-reform exposure to minors is strongly correlated with firm size, and larger manufacturing firms followed different post-2009 trajectories through the crisis and recovery, omitting size-specific time paths allows the binary treatment to pick up scale-related import dynamics. The positive robot coefficients are concentrated in the binary and percentile codings and disappear under the continuous-share treatment, which further suggests that they reflect the size gradient rather than a reform-specific response. This comparison illustrates why the size-adjusted specification is the chapter's baseline.

The remaining appendix sensitivities narrow the interpretation without changing which outcomes are positive and which are null. Tighter binary exposure cutoffs (5% and 10% of pre-reform minor share, Table B16) preserve a positive but smaller ICT coefficient and do not uncover a positive import response. Adding baseline firm age and workforce average age to the control set (Table B8) attenuates ICT modestly. The single-cohort imputation ATT (Table B17) leaves the annual importer at zero and ICT at 0.0143–0.0183, close to the baseline. The demanding checks are less favourable. The trend-augmented specification, which allows more-exposed firms to follow a differential linear trend, moves ICT to -0.0031 in Panel A and 0.0086 in Panel B – essentially eliminating the result under the earlier timing window and halving it under the later one. The matched common-support design (Table B19) produces a similar pattern: ICT

falls to 0.0023 and 0.0056 in the matched TWFE and becomes -0.0029 and 0.0085 with added trends. B.3 adds the same caution from another angle: the age 19–21 and age 20–22 placebo exposures yield ICT coefficients around 0.0240 and 0.0234 in Panel A, and 0.0240 and 0.0244 in Panel B, all larger than the baseline binary estimates, while the fake-timing grid produces small positive pseudo-effects in some pre-reform years. These exercises do not overturn the ICT result mechanically, but they do show that it is bounded and sensitive to comparability. The import margins remain null throughout.

The overall picture is stable across these exercises. The reform is associated with a strong first-stage reduction in underage employment. Complementary ICT employment is the clearest downstream response, but it is attenuated by tighter comparability designs and its specificity is weakened by the placebo evidence. The occupation-based proxy points toward internal workforce recomposition but is less precisely estimated. Annual robot-import indicators remain close to zero under every specification. The chapter therefore documents differential post-reform firm adjustment, concentrated in workforce composition rather than equipment imports, and with the causal interpretation constrained by the design limitations discussed above.

Appendix B.2 reports additional robustness exercises. Taken together, they preserve the same broad ordering of outcomes. Across alternative functional forms and windows, the ICT margin remains the clearest positive response, while the strict Tier 1 and Tier 2 robot outcomes remain weak once the baseline size-adjusted design is imposed. The same appendix also compares the binary TWFE estimates with a single-cohort imputation ATT, a trend-augmented TWFE sensitivity, a matched common-support exercise, and a short transformation robustness check that replaces $\log(1+Y)$ with $\operatorname{asinh}(Y)$ for the cumulative import outcomes. These alternatives are generally more conservative than the baseline TWFE estimates, especially once differential trends or tighter common-support restrictions are imposed. An additional clustering comparison is also reported but does not change the broad ordering of outcomes. Appendix B.3 reports the broader placebo battery, which is best read as cautionary: negative-control exposure placebos are mostly null for the robot outcomes but remain clearly positive for ICT, while the fake-timing grid produces negative pseudo-effects for the annual importer and small positive pseudo-effects for ICT in some pre-reform years. Having established the average response, the next subsection asks where in the firm distribution that response is concentrated.

3.6.4 Heterogeneity by firm size and labour intensity

The average effects in Table 16 are not shared equally across firms. Table 17 re-estimates the binary-exposure specification within pre-reform size quartiles and labour-intensity subgroups, and reports the triple-difference from Equation 8. We treat these results as descriptive concentration patterns rather than as a second source of identification.

The clearest heterogeneity again appears on ICT. In the size-quartile estimates, the ICT coefficient is essentially zero in the bottom quartile in both panels, modest and imprecise through the middle of the distribution, and clearly positive only in the top quartile, at 0.0180 in Panel A and 0.0242 in Panel B. The

Table 16: Firm-level DiD estimates for four annual adjustment outcomes

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Panel A: Post 2009				
Minors × Post 2009	-0.0002	-0.0000	0.0082	0.0116***
	(0.001)	(0.001)	(0.005)	(0.004)
Share minors × Post 2009	-0.0019	-0.0025	0.0490**	0.0337**
	(0.002)	(0.003)	(0.022)	(0.017)
Percentile share minors × Post 2009	-0.0001	-0.0002	0.0041**	0.0037***
	(0.000)	(0.000)	(0.002)	(0.001)
Observations	428,076	428,076	428,076	428,076
Panel B: Post 2012				
Minors × Post 2012	-0.0003	0.0003	0.0127**	0.0163***
	(0.001)	(0.001)	(0.006)	(0.004)
Share minors × Post 2012	-0.0035*	-0.0031	0.0535**	0.0435**
	(0.002)	(0.003)	(0.023)	(0.019)
Percentile share minors × Post 2012	-0.0002	-0.0001	0.0051***	0.0047***
	(0.000)	(0.000)	(0.002)	(0.001)
Observations	353,203	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual Tier 1 importer, annual Tier 2 importer, occupation proxy (flexible), and ICT worker. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the three treatment codings used in the main text: the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. "Minors" is a dummy equal to 1 if the firm employed workers under 18 during 2006–2008; "Share of minors" is the corresponding pre-reform workforce share; and "Percentile share minors" is the discretized version of that share used in the manuscript. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies. Appendix Table B4 reports the related five-outcome block that adds cumulative import outcomes.
Source: Authors' calculations using QP and CI data, 2002–2019.

triple-difference sharpens the same message: the DiD coefficient for smaller exposed firms ($\hat{\beta}_1$) is close to zero on ICT, at 0.0002 in Panel A and 0.0013 in Panel B, while the incremental large-firm interaction ($\hat{\beta}_2$) is 0.0133 and 0.0172. The pattern is already present in Panel A and is somewhat sharper in Panel B, which is consistent with delayed adjustment but also reflects conditioning on survival through 2012 rather than a clean timing comparison. The flexible occupation proxy points in the same direction, but less cleanly. It is positive in the upper quartiles, reaches 0.0210 in the top quartile of Panel B, and its large-firm incremental effect is 0.0131 in Panel A and 0.0161 in Panel B. This is consistent with a more visible internal reorganization margin among larger firms, but it remains suggestive rather than decisive.

The import outcomes do not display a comparable size gradient. Tier 1 remains close to zero in every subgroup, and the Tier 2 triple-difference yields a small positive large-firm increment only in Panel B (0.0032), without a clean counterpart in Panel A. The labour-intensity split is weaker than the size split: ICT stays positive in both subgroups, while the occupation proxy is clearly positive only in the low-labour-cost-share subgroup of Panel B.

The concentration of the ICT response among larger firms is consistent with greater organisational capacity to adjust workforce composition after a labour-supply shock. Larger firms are more likely to absorb the loss of cheap young workers by hiring complementary technical staff. But this pattern is not additional causal validation. The binary exposure indicator partly captures a firm-size gradient, so the heterogeneity results may also reflect broader size-related differences rather than a reform-specific mechanism. We therefore treat them as descriptive refinements of the average effect.

3.6.5 Firm exit and attrition robustness

This subsection asks whether differential exit or attrition among more exposed firms could drive the main results. Table 19 reports firm fixed-effects linear probability estimates of exit on the SCIE-linked sample under the three main treatment codings. Under the binary exposure, the coefficients are -0.0023 in Panel A and -0.0020 in Panel B, both statistically insignificant. The continuous-share and percentile-share codings tell the same story: no treatment coding yields a stable positive exit coefficient across the two panels. The data do not support a reform-induced wave of exit among more exposed firms. Appendix Table B3 reports a pooled logit benchmark without firm fixed effects as a descriptive comparison.

Table 20 then asks the more direct question: if we restrict the sample to firms that survive the entire panel, eliminating differential attrition mechanically, do the main results change? The table re-estimates the main binary-exposure TWFE specification for the same four outcomes as Table 16 under two balanced-panel restrictions: firms observed for at least 12 of the 18 panel years (roughly 20,000 firms) and firms observed for all 18 years (roughly 8,000 firms). The annual Tier 1 importer remains null under both restrictions in both panels. The occupation proxy retains its positive sign and becomes marginally significant in the fully balanced Panel B sample (0.0167, $p = 0.07$). ICT is the outcome that changes most: it rises from 0.0116 in the main sample to 0.0239 in the fully balanced Panel A sample ($p < 0.001$), and from 0.0163 to 0.0259 in Panel B ($p < 0.001$). The cumulative Tier 1 import coefficient

becomes more negative in the balanced samples, consistent with the earlier finding that it is sensitive to sample composition.

The balanced-panel results have two implications. First, differential attrition is not inflating the ICT result. If it were, restricting the sample to survivors should weaken the coefficient, not strengthen it. The fact that ICT roughly doubles in the fully balanced sample suggests that firms that exit the sample are firms with weaker ICT responses, so the main-sample estimate is, if anything, attenuated by attrition. Second, the concentration of the ICT effect among long-lived firms is consistent with the heterogeneity findings in Section 3.6.4: larger, more persistent firms are those with the organisational capacity to adjust their workforce composition. This does not resolve the underlying selection concern - long-lived firms differ from short-lived firms on many dimensions - but it does rule out the specific concern that differential exit is driving the positive ICT coefficient. The sector–region cell-level evidence reported in Appendix Table B21 reinforces this reading by showing that the firm-level effects do not aggregate into a broader market-level response.

3.7 Conclusion

Portugal's 2009 compulsory schooling reform raised the minimum school-leaving age from 15 to 18 and reduced the supply of underage workers to manufacturing firms. This chapter asks whether firms more exposed to that labour-supply shock adjusted differently after the reform. Using linked employer–employee and trade data, we estimate a common-timing Difference-in-Differences design over 2002–2019.

The first stage is strong. The share of minors fell by 3.15 percentage points in Panel A and 2.23 in Panel B; the probability of employing any minor fell by 28.6 and 24.0 percentage points. These are large relative effects – the share of minors more than halved in exposed firms – but the absolute level of the affected input was small: minors represented roughly 3–5% of the workforce in treated firms before the reform. Any downstream adjustment should be interpreted against that scale.

Among the three margins we examine, complementary ICT employment shows the clearest differential post-reform response in the baseline design. Under the binary exposure, more exposed firms are 1.2 percentage points more likely to employ ICT workers in Panel A and 1.6 percentage points in Panel B, and the result is positive under all three treatment codings. An occupation-based proxy for automation-related workforce presence is also positive, especially in Panel B and under the continuous-share treatment, but it is less precisely estimated and we treat it as suggestive evidence of internal recomposition. Annual robot-import indicators are close to zero in both panels and under every specification. The import margins carry an additional measurement limitation: robot purchases are extremely rare events at the firm level, so the power to detect even a genuine effect on that margin is limited. These three margins capture different forms of adjustment – equipment acquisition, internal occupational restructuring, and complementary technical employment – and the evidence points much more clearly to the last of these than to the first.

The main interpretive limits are clear. More- and less-exposed firms are not perfectly comparable,

and the reform overlaps with the sovereign-debt crisis. The ICT result is also weakened by the diagnostic evidence. Age-based exposure placebos – which replace minors with young workers aged 19–21 or 20–22 who were not directly targeted by the reform – produce ICT coefficients of 0.0234–0.0244, larger than the baseline 0.0116 and 0.0163, and the fake-timing exercises produce small positive pseudo-effects in some pre-reform years. Trend-augmented specifications reduce the ICT coefficient to -0.0031 in Panel A and 0.0086 in Panel B. Matched common-support designs reduce it to 0.0023 and 0.0056, and to -0.0029 and 0.0085 once matched trends are added. The ICT result therefore does not survive the most demanding comparability checks fully intact. At the same time, it is not driven by differential attrition: a balanced-panel restriction shows that the coefficient roughly doubles among firms observed for all 18 panel years, consistent with the heterogeneity finding that the response is concentrated in larger, more persistent firms. Exit results do not support a cleansing channel, and sector–region cell-level regressions confirm that the within-firm effects do not aggregate into a broader market-level response (Appendix Table B21).

The directed technical change literature predicts that labour scarcity encourages labour-saving innovation. The evidence here is only partially consistent with that prediction. The reform tightened a specific labour input, and more exposed firms adjusted through complementary technical employment and, more tentatively, through occupational recomposition. But the response appears in workforce composition, not in equipment acquisition. That pattern is consistent with the scale of the shock: losing 3–5% of the workforce in the form of the cheapest and least-skilled workers may not generate enough pressure to trigger capital investment in robots, while it may still be enough to shift hiring toward complementary technical occupations, especially in larger firms with greater organisational capacity.

As a thesis chapter, the contribution is to establish a strong labour-supply first stage and to document a differential post-reform pattern of firm adjustment while keeping the main interpretive limits explicit. The clearest finding is higher complementary ICT employment in more exposed firms, but that pattern is not reform-specific in a clean sense. The occupation-based proxy is suggestive rather than decisive. We do not find robust evidence that the reform increased robot importing. Whether the ICT pattern reflects the reform itself or a broader process of workforce upgrading among firms that relied more on young labour remains an open question.

Table 17: Heterogeneity by firm size and labour intensity in four annual adjustment outcomes

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Panel A: Post 2009				
Size quartiles				
Size Q1	-0.0001**	-0.0001	-0.0010	-0.0006
	(0.000)	(0.000)	(0.003)	(0.002)
Size Q2	0.0002	-0.0007	-0.0008	0.0052*
	(0.000)	(0.001)	(0.004)	(0.003)
Size Q3	0.0000	0.0013	0.0085	0.0012
	(0.000)	(0.001)	(0.006)	(0.003)
Size Q4	-0.0001	-0.0004	0.0134	0.0180***
	(0.001)	(0.002)	(0.009)	(0.007)
Labour intensity				
Low labour-cost share	-0.0002	0.0006	0.0139	0.0109*
	(0.001)	(0.002)	(0.009)	(0.006)
High labour-cost share	-0.0003	-0.0013	0.0014	0.0135***
	(0.001)	(0.001)	(0.007)	(0.005)
Triple difference				
DiD for small firms	-0.0005	-0.0016**	-0.0029	0.0002
	(0.000)	(0.001)	(0.003)	(0.003)
Large-firm incremental effect	0.0003	0.0017	0.0131*	0.0133***
	(0.001)	(0.001)	(0.007)	(0.005)
Firm FE	Yes	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes	Yes

Table 18: Heterogeneity by firm size and labour intensity in four annual adjustment outcomes (continued)

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Panel B: Post 2012				
Size quartiles				
Size Q1	-0.0001	0.0000	-0.0022	0.0002
	(0.000)	(0.000)	(0.003)	(0.003)
Size Q2	-0.0007	-0.0014*	0.0036	0.0070
	(0.001)	(0.001)	(0.004)	(0.005)
Size Q3	-0.0001	0.0020*	0.0079	0.0035
	(0.000)	(0.001)	(0.007)	(0.004)
Size Q4	-0.0001	-0.0003	0.0210**	0.0242***
	(0.002)	(0.002)	(0.010)	(0.007)
Labour intensity				
Low labour-cost share	0.0002	0.0012	0.0186**	0.0176***
	(0.002)	(0.002)	(0.009)	(0.007)
High labour-cost share	-0.0012	-0.0014	0.0060	0.0160***
	(0.001)	(0.001)	(0.008)	(0.005)
Triple difference				
DiD for small firms	-0.0014**	-0.0026***	-0.0011	0.0013
	(0.001)	(0.001)	(0.003)	(0.004)
Large-firm incremental effect	0.0012	0.0032**	0.0161**	0.0172***
	(0.001)	(0.002)	(0.007)	(0.006)
Firm FE	Yes	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report the same four annual outcomes as Table 16. The size-quartile and labour-intensity rows report the coefficient on the binary exposure interacted with the relevant post indicator within the stated subsample. The triple-difference rows report $\hat{\beta}_1$ and $\hat{\beta}_2$ from Equation 8, where $\hat{\beta}_1$ is the DiD coefficient for smaller firms and $\hat{\beta}_2$ is the incremental effect for larger firms defined as above-median baseline employment. Labour-intensity groups are defined using the median of the baseline labour-cost share from SCIE. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies. The first page reports Panel A, which includes firms that existed in 2009 and uses the Post-2009 coding; the continued page reports Panel B, which includes firms that existed in 2012 and uses the Post-2012 coding.
Source: Authors' calculations using QP, SCIE, and CI data, 2002–2019.

Table 19: Firm-level DiD estimates for firm exit in the SCIE-linked sample

	Market exit
Panel A: Post 2009	
Minors x Post 2009	-0.0023
	(0.002)
Share minors x Post 2009	-0.0421***
	(0.014)
Percentile share minors x Post 2009	-0.0007
	(0.001)
Observations	279,098
Panel B: Post 2012	
Minors x Post 2012	-0.0020
	(0.001)
Share minors x Post 2012	-0.0202
	(0.013)
Percentile share minors x Post 2012	-0.0007
	(0.000)
Observations	243,376
Firm FE	Yes
Controls × Year	Yes

Notes: Difference-in-differences linear probability model corresponding to Equation 5, estimated on the SCIE-linked exit sample with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. The outcome equals 1 in the firm's last observed year of SCIE activity. Rows report the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using SCIE and QP data, 2006–2019.

Table 20: Balanced-panel robustness: main adjustment outcomes under survivor restrictions

	Tier 1 importer	Cum. Tier 1 imports (log)	Occ. flexible proxy	ICT worker
Panel A				
Main sample	-0.0002	-0.0400*	0.0082	0.0116***
	(0.0007)	(0.0205)	(0.0055)	(0.0038)
N = 428,076 [34,297 firms]				
Balanced 12+ years	0.0000	-0.0452*	0.0082	0.0122***
	(0.0009)	(0.0241)	(0.0063)	(0.0044)
N = 321,687 [20,081 firms]				
Fully balanced	-0.0001	-0.0821**	0.0105	0.0239***
	(0.0014)	(0.0381)	(0.0095)	(0.0063)
N = 143,802 [7,989 firms]				
Panel B				
Main sample	-0.0003	-0.0421*	0.0127**	0.0163***
	(0.0009)	(0.0235)	(0.0058)	(0.0042)
N = 353,203 [23,740 firms]				
Balanced 12+ years	-0.0003	-0.0420*	0.0133**	0.0164***
	(0.0009)	(0.0247)	(0.0060)	(0.0044)
N = 315,136 [19,595 firms]				
Fully balanced	-0.0009	-0.0657*	0.0167*	0.0259***
	(0.0016)	(0.0393)	(0.0091)	(0.0064)
N = 143,802 [7,989 firms]				
Firm FE	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes
Region × Year FE	Yes	Yes	Yes	Yes
Baseline controls × year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, using the binary exposure and the same firm, sector × year, and region × year fixed effects as Table 16. "Main sample" reproduces the baseline estimates. "Balanced 12+ years" restricts to firms observed for at least 12 of the 18 panel years. "Fully balanced" restricts to firms observed in all 18 years (2002–2019). Standard errors are clustered at the firm level. Observation counts and the number of firms are shown for each restriction. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding.
Source: Authors' calculations using QP and CI data, 2002–2019.

4

Minimum wage, robot adoption, and worker adjustment: Evidence from Portugal

4.1 Introduction

Minimum Wage (MW) policies are a central instrument for raising the earnings floor and reducing wage inequality.[1] At the same time, production processes are undergoing a structural transformation driven by the diffusion of automation technologies, including industrial robots. These developments generate a fundamental tension: by increasing the relative cost of low-wage labour, minimum wage increases may strengthen firms' incentives to substitute capital for labour, potentially accelerating automation and reshaping the allocation of tasks and jobs within firms.

We study whether minimum wage increases induce firms to automate and how the resulting adjustment costs are distributed across workers. Using Portuguese manufacturing firms from 2010 to 2021 and combining administrative firm and matched employer-employee data with international trade records of industrial robot imports (which we use to proxy robot purchases), we exploit predetermined firm-level exposure to statutory minimum wage increases to trace dynamic responses over time. We find that minimum-wage-induced cost shocks raise the likelihood of robot adoption, but the response materialises only after a delay of three to five years. Conditional on adoption, automation reshapes worker outcomes within firms: among workers who remain in the same firm and occupation, wage growth is less favourable for those in more automatable tasks, while job separations rise disproportionately for workers on fixed-term contracts and observed internal reallocation for that group is limited.

Despite these conceptual links, the empirical understanding of how labour cost shocks trigger automation remains underdeveloped. The extensive literature on minimum wages has primarily focused on short-run disemployment effects, often neglecting the medium-run margin of technological substitution. Conversely, studies on the determinants of robot adoption often treat automation as an exogenous shock driven by falling global technology prices, paying less attention to the role of domestic policy in accelerating diffusion. Establishing causality is further complicated by positive selection: high-productivity firms are

[1] This chapter is co-authored with Hélder Costa.

both more likely to pay higher wages and more likely to adopt robots, so separating the effect of a policy shock from unobserved firm heterogeneity requires a rigorous identification strategy.

Accordingly, we ask two questions. First, do increases in the statutory minimum wage induce firms to adopt industrial robots, and over what horizon do these responses materialise? Second, conditional on adoption, how does automation reshape workers' outcomes within firms – through wages, separations, and occupational reallocation – and how are these margins distributed across workers in more versus less automatable tasks, particularly among fixed-term workers in a segmented labour market?

Theoretical predictions regarding these mechanisms are ambiguous. Task-based frameworks emphasise that automation displaces labour from tasks in which it previously held a comparative advantage, while the creation of new tasks can reinstate labour demand (Acemoglu & Restrepo, 2018b, 2022b). In this setting, increases in labour costs raise the relative price of low-wage labour, potentially accelerating technology adoption (Dai & Qiu, 2024; Fan et al., 2021). At the same time, robot adoption entails substantial fixed costs, organisational restructuring, and financing constraints. As a result, minimum wage hikes may initially reduce firms' capacity to invest – particularly among less productive firms – delaying or discouraging automation. These opposing forces suggest that firms' responses to minimum wage increases are heterogeneous and dynamic, and may unfold only over the medium run.

Portugal provides a particularly suitable setting to study these dynamics due to the magnitude of the policy shock and the structural rigidity of its labour market. Between 2014 and 2021, the country implemented a sequence of large, nationally uniform minimum wage increases, raising it by around 30% in real terms, which coincided with a period of rapid industrial automation. Crucially, the Portuguese labour market is characterised by a high degree of segmentation, combining strict employment protection for permanent workers with comparatively low protection for those on fixed-term contracts. This institutional context allows us to examine not only whether firms automate in response to cost shocks, but also whether the resulting adjustment is concentrated among workers on fixed-term contracts, as a dual-labour-market perspective would suggest.

Empirically, our analysis proceeds in two stages. First, we estimate the dynamic response of robot adoption to minimum wage exposure using a predetermined measure that isolates variation driven exclusively by statutory minimum wage changes – holding baseline wage structures fixed so that changes reflect policy-implied increases in labour costs rather than endogenous wage adjustments or compositional shifts. This design allows us to trace medium-run adjustment while mitigating concerns about reverse causality and time-varying unobservables correlated with both wages and automation. Second, exploiting within-firm variation in workers' pre-adoption exposure to automation (measured before adoption and then frozen), we study heterogeneous effects on wages, job separations, and occupational reallocation within adopting firms. Complementary descriptive firm-level trajectories around adoption are reported separately and used as selection diagnostics rather than as causal evidence on firm performance.

The paper makes two main contributions. First, we bring a policy-driven source of labour cost variation to the study of automation by providing plausibly causal evidence on the medium-run technology margin of minimum wage increases, documenting sizeable adjustment lags consistent with lumpy investment and

planning frictions. Second, we connect firms' adoption decisions to within-firm distributional outcomes by showing how automation affects workers differently depending on pre-adoption task exposure, with the clearest adjustment concentrated among fixed-term workers. The pattern of concentrated separations among fixed-term workers and limited observed internal reallocation for that group is consistent with a dualised adjustment process in a segmented labour market (Boeri & Garibaldi, 2007; Centeno & Novo, 2012).

Our findings have direct implications for policy in an era of rapid technological change. We show that while minimum wages effectively raise the earnings floor, they also catalyse a medium-run substitution towards capital that reshapes labour demand within firms. Crucially, the burden of this adjustment is uneven: displacement concentrates on fixed-term contracts, consistent with a dual labour market in which stronger protection for insiders may shift adjustment towards more precarious workers. At the same time, wage estimates among stayers indicate that even retained workers in more automatable tasks experience relative wage penalties. This highlights a critical policy trade-off: in labour markets with strong employment protection for permanent workers, the costs of automation may concentrate on the most precarious segments of the workforce.

The remainder of the paper is organised as follows. Section 4.2 reviews the related literature. Section 4.3 describes the data and variable construction. Section 4.4 outlines the empirical strategy. Section 4.5 presents the results. Section 4.6 concludes.

4.2 Literature review

This paper contributes to three closely related strands of literature: the employment effects of minimum wages, the determinants of technology adoption, and the firm- and worker-level consequences of automation. While standard models predict that higher labour costs induce capital–labour substitution, direct empirical evidence linking minimum-wage-induced labour-cost shocks to industrial-robot adoption – and to the subsequent medium-run adjustment dynamics – remains comparatively scarce.

Moreover, much of the existing literature abstracts from labour market institutions, such as employment protection legislation, that can shape both the timing of adoption and the distribution of adjustment costs across workers.

4.2.1 Minimum wage, labour costs, and robot adoption

Our analysis builds on the task-based framework of technological change developed by Acemoglu and Restrepo (2018b), in which production consists of a set of tasks that can be performed by either labour or capital. Automation expands the range of tasks executed by capital, displacing workers from routine and codifiable activities while simultaneously creating new tasks in which labour retains a comparative advantage (Acemoglu & Restrepo, 2018b, 2019; Autor, 2015). The routine-task framework introduced by Autor et al. (2003) and the evidence on job polarisation documented by Autor and Dorn (2013) provide the

empirical foundations for this view. A central implication of these models is that increases in labour costs strengthen incentives to adopt labour-saving technologies. This mechanism is supported by evidence that changes in relative factor prices – whether driven by immigration-induced labour supply shifts (Lewis, 2011), aging demographics reducing the supply of middle-aged production workers (Acemoglu & Restrepo, 2022a) or by declining robot prices (Cheng et al., 2019) – alter the capital–labour mix, and it is increasingly examined in the context of minimum wage policies.

A growing empirical literature provides evidence that mandated wage floors can accelerate capital–labour substitution. Studies using Chinese data offer particularly clear support: Hau et al. (2020) show that minimum-wage-induced cost shocks raise capital intensity and accelerate total factor productivity growth, with the strongest effects among initially less productive private and foreign-owned firms and little evidence of a surge among state-owned enterprises. Fan et al. (2021) find that higher minimum wages raise the likelihood of robot adoption, especially among more productive firms, privately owned firms, and those located in coastal provinces. In high-income economies, Lordan and Neumark (2018) document that minimum wage increases reduce the share of automatable employment held by low-skilled workers, with especially pronounced effects in manufacturing. Along similar lines, Dai and Qiu (2024) show that U.S. establishments exposed to minimum wage hikes expand investment in information technology. In the most direct U.S. evidence on minimum wages and industrial robot adoption, Brynjolfsson et al. (2026) link firm-level robot import records for single-unit manufacturing plants from the Longitudinal Firm Trade Transactions Database to Census microdata for 1992–2021 and exploit geographic discontinuities at state borders to identify causal effects; they find that a 10% increase in the minimum wage raises the probability of robot adoption by approximately 8% relative to the sample mean. Relatedly, J. Clemens et al. (2021) show that minimum wage increases induce skill upgrading: workers employed in low-wage occupations become older and more likely to hold a high school diploma, and firms raise educational requirements in job postings, consistent with labour–labour substitution from lower- to higher-skilled workers. Beyond manufacturing, Aaronson and Phelan (2019) show that minimum wage hikes disproportionately reduce employment in low-wage occupations intensive in routine cognitive tasks, consistent with technological substitution of such tasks.

Minimum wage increases may also induce adjustment along margins other than technology adoption. Firms can respond through changes in prices, profits, and product-market strategies. Harasztosi and Lindner (2019) show that Hungarian firms absorb minimum wage hikes through a combination of higher output prices, lower profit margins, and increased productivity rather than through employment cuts. In a different institutional setting, Dustmann et al. (2022) find that German minimum wage increases trigger reallocation of workers from smaller to larger and more productive firms, with compositional upgrading rather than aggregate disemployment.[2]

Most causal evidence on automation adoption comes from a small set of advanced economies – particularly the United States, France, Germany and Spain – and therefore provides limited guidance for

[2] At a more aggregate level, Cengiz et al. (2019) show that missing jobs below the new minimum wage are largely matched by excess jobs at and slightly above it, leaving low-wage employment essentially unchanged while raising earnings.

institutional environments with stronger employment protection and segmented labour markets. This limitation is particularly relevant because adoption incentives depend on the adjustment margins available to firms. For Portugal, available evidence shows that robot adoption is comparatively low but concentrated among more productive and export-oriented firms (Amador, 2024), and that the adoption of digital and Industry 4.0 technologies is subject to strong positive selection by firm size and pre-existing digital capabilities (Alexandre et al., 2026; Barbosa, 2024; Barbosa & Faria, 2022); at the occupational level, Casas et al. (2025) document substantial heterogeneity in exposure to automation and AI across the Portuguese workforce. Using Portuguese trade data, Bastos et al. (2023) show that increased product-market competition generates heterogeneous automation responses: on average, greater competitive pressure reduces automation investment, but the most productive firms in highly automating industries respond by increasing it. More recently, Tiedtke (2025) uses Portuguese administrative data and a difference-in-differences event study to show that automation creates jobs in small firms but leads to job losses in larger ones, with effects that favour low-educated and routine blue-collar workers over Science, Technology, Engineering and Mathematics (STEM) professionals. However, these studies focus on diffusion patterns, market competition, and employment composition rather than on labour-cost shocks as a driver of adoption. As a result, there is no causal evidence on whether binding wage floors induce automation in a Southern European dual labour market. The Portuguese minimum-wage increases therefore provide a setting to connect the literature on policy-induced labour demand adjustment with endogenous technology adoption, testing whether firms substitute towards robots when adjustment through internal task reallocation is institutionally constrained.

Overall, while this literature identifies several channels through which minimum wage increases may affect automation decisions, much less is known about the timing of these responses and about how labour market rigidities shape their transmission. Our analysis speaks directly to this gap. We measure firm-level exposure to the minimum wage using the PRCI index of Alexandre et al. (2022), which captures the proportional wage-bill increase required to bring all sub-minimum workers up to the new floor. Combining this continuous treatment intensity with the heterogeneity-robust difference-in-differences estimator of Callaway and Sant'Anna (2021), we document delayed adoption responses consistent with models of lumpy capital adjustment and fixed investment costs (e.g., Cooper & Haltiwanger, 2006), and highlight the mediating role of employment protection in dual labour markets.

4.2.2 Automation and firm-level outcomes

Theoretical models of automation predict two opposing forces at the firm level: a displacement effect that substitutes capital for labour, and a productivity effect that lowers marginal costs and induces output expansion (Acemoglu & Restrepo, 2018b; Acemoglu et al., 2020, 2023). The net employment impact depends on whether the scale effect generated by higher competitiveness outweighs direct task substitution.

Empirical evidence points to sizable productivity gains from robotization, but employment effects depend on the margin and level of analysis. At the industry/country level, robot penetration is associated with faster labour productivity growth and higher TFP, alongside lower output prices (Graetz & Michaels, 2018). Using French firm-level data on robot imports, Bonfiglioli et al. (2024) document strong selection into importing and find that exogenous exposure to robots raises productivity and the share of high-skill employment, while effects on sales and total employment are more muted. Using French administrative and trade data, Acemoglu et al. (2020) find that robot-adopting firms expand value added and productivity and may increase their own employment, but adoption also reduces competitors' employment, implying negative net effects at the industry level. Related evidence using broader measures of automation for France is provided by Aghion et al. (2020). In Japan, Adachi et al. (2024) find that declines in robot prices increased employment via higher productivity and scale at the industry and local level. For the Netherlands, Acemoglu et al. (2023) conclude that workers directly affected by robots experience lower earnings and reduced employment rates after adoption.

A key challenge in this literature is disentangling causal effects from positive selection into adoption. Koch et al. (2021) show that more productive Spanish firms are more likely to adopt robots, but also find substantial causal treatment effects: output gains in the vicinity of 20–25% within four years, net job creation of around 10%, and a reduction in the labour cost share by about 5–7 percentage points. Importantly, they also document negative reallocation effects on non-adopters, who lose market share and face higher exit rates – a pattern consistent with the competitive dynamics identified by Acemoglu et al. (2020) in French data. Related evidence underscores strong heterogeneity and selection into adoption across settings (Bessen et al., 2020; Freeman et al., 2025). Beyond scale, Dinlersoz and Wolf (2024) document declining labour shares following automation, while Dixon et al. (2021) highlight that productivity gains may also arise from improved quality consistency and reduced monitoring costs, reshaping occupational structures without necessarily reducing total employment.

Taken together, this literature suggests that robot adopters tend to expand and gain market share, reflecting both pre-existing advantages and genuine causal gains from automation (Koch et al., 2021). Whether expansion translates into changes in labour shares or workforce composition appears to depend on institutional context, motivating our separation of descriptive firm-level patterns (Appendix C.3) from causal worker-level estimates in the main analysis.

4.2.3 Worker-level effects and task reallocation

While the firm-level evidence reviewed above captures aggregate outcomes for adopters, it does not reveal how adjustment costs are distributed across individual workers. The effects of automation on individual workers are highly heterogeneous and closely linked to task content and occupational exposure. Seminal work by Acemoglu and Restrepo (2020) shows that greater robot exposure reduces employment and wages in local labour markets, with adverse effects concentrated in routine-intensive occupations. Complementary evidence at the commuting-zone level is provided by Autor et al. (2015), who show that

routine-biased technological change reshapes local labour markets, a pattern closely related to the routine task intensity measures used in this paper.

Evidence from regional labour markets and firm-level panels highlights heterogeneous adjustment channels. Using German administrative data and a shift-share design, Dauth et al. (2021) show that robot exposure reduces manufacturing employment but is fully offset by job growth in services, with the incidence of displacement concentrated on young labour market entrants; incumbent workers, by contrast, see increased plant tenure and transition into higher-quality tasks within their original firms. Using matched employer–employee data for the Netherlands, Bessen et al. (2025) show that automation increases separation and non-employment durations and leads to sizeable cumulative earnings losses driven by fewer days worked. This evidence highlights that even when firms expand, workforce composition may shift against routine-task workers (Autor et al., 2003). Humlum (2021) combines event studies with a structural occupational-choice model for Denmark, emphasizing distributional effects and substantial barriers to reallocation, especially for older workers. At the same time, Freeman et al. (2025) find that robot adoption in China can raise productivity and market share and is accompanied by higher employment and wages within adopting firms, indicating that displacement and reinstatement effects may coexist.

A critical yet underexplored dimension of these dynamics is their interaction with labour market institutions, especially in dual labour markets characterised by a divide between permanent and temporary contracts. Such settings can concentrate adjustment on the temporary margin (Boeri & Garibaldi, 2007; Centeno & Novo, 2012). Blanas et al. (2019) provide cross-country evidence that robots shift labour demand away from low-skill and young workers and towards high-skill workers, with particularly strong effects in manufacturing. In rigid labour markets, firms may therefore absorb shocks primarily through non-renewal of temporary contracts, reinforcing the dualisation of employment risks.

These findings underscore that the distributional consequences of automation depend not only on the balance between displacement and the reinstatement of labour through new tasks (Acemoglu & Restrepo, 2022b), or on occupational exposure to automation (Montobbio et al., 2024), but also on the institutional rules governing hiring and separations.[3]

By explicitly analysing contract heterogeneity, our study contributes to a more complete understanding of how technological change affects workers in regulated labour markets.

4.3 Data, variables description, and summary statistics

This section describes the datasets used in this paper, defines the key variables, and provides summary statistics.

[3]Complementary patent–task text-matching approaches provide alternative measures of occupational exposure. Webb (2020) construct exposure scores from verb–object overlap between patent titles and Occupational Information Network (O*NET) task descriptions; as a robustness check, we use the labour-saving robotics exposure index of Montobbio et al. (2024), which matches Cooperative Patent Classification (CPC) code definitions from labour-saving robot patents to O*NET tasks (Appendix C.4).

4.3.1 Data

This study leverages data from three linked datasets for Portugal covering the years 2010–2021. The first is the Integrated Business Accounts System, SCIE (*Sistema de Contas Integradas das Empresas*, INE, 2021b) from Statistics Portugal's Registry of Statistical Units. The SCIE database is annual and contains firm-level administrative data on balance sheets and other accounting information for the population of non-financial private-sector firms. In 2021, the SCIE database included 468,746 firms. From SCIE we collect annual firm-level data on value added, turnover, employment, industry classification, total costs, and exports - as detailed in Table 21 and Table 22.

The second dataset is the linked employer–employee Personnel Records database QP, (*Quadros de Pessoal*, INE, 2021c), collected by the Portuguese Ministry of Labour, Solidarity and Social Security (MTSSS) through a mandatory annual survey. QP is an annual linked employer–employee dataset covering all private-sector firms in Portugal with at least one wage earner, and all their workers, together with firm-level information (including sales, location, and economic activity). In 2021, the QP database contained 290,409 firms and 3,132,695 workers. From this dataset, we collect worker-level information, such as occupation, education, age, wages, and hours worked, as well as firm age (see Table 21 and Table 22).

The third dataset is the monthly CI (*Comércio Internacional*, INE, Statistics Portugal, 2025), which provides firm-level information on export and import values, as well as detailed product codes that we use to identify the adoption of automation technologies via capital-goods purchases (in 2021 it comprised 158,717 firms). The CI dataset covers all trade transactions involving Portuguese firms, distinguishing between intra-European Union (EU) and extra-EU trade. Trade with non-EU countries is comprehensively recorded using customs administrative data. Intra-EU trade is collected through the Intrastat survey, which applies a reporting threshold that is adjusted annually to cover at least 93–95% of total arrivals and 97% of total dispatches, as mandated by EU regulations. As noted by Acemoglu et al. (2023), who use comparable data for the Netherlands, this reporting threshold has a negligible impact on data completeness for major capital investments such as industrial robots. Firms that do not appear in the international trade have their import values set to 0.

The three datasets are linked through an anonymized firm identifier, allowing us to analyze the relationship between firm- and worker-level characteristics and international trade.

4.3.2 Sample and variables description

Our analysis focuses on the sample of manufacturing firms. Only firms for which information is available in all datasets are considered. We follow the OECD Structural Analysis Database – A7 industry classification (STAN A7) classification, which define the manufacturing sector as the 2-digit Statistical Classification of Economic Activities in the European Community (Revision 2) (NACE Rev. 2) industries between codes 10 and 33, inclusive. Our final sample consists of approximately 23,000 firms per year, with information on about 500,000 workers annually. In 2010, a total of 90 firms possessed a positive stock of robots, a

figure that rose to 377 by 2021 (see Figure 6). Although robot adopting firms constitute a small fraction of the total firm population in the manufacturing sector, they account for a disproportionately large share of economic activity, capturing approximately 19% of total value added and 14% of employment in 2021.

In this study, robotics refers exclusively to the use of industrial robots.[4] To measure robot adoption, this study relies on international trade data that tracks imports of capital goods. Following the methodology of Acemoglu et al. (2020) and Acemoglu et al. (2023), robot adoption is proxied using trade flows classified under the CN code 847950, which specifically identifies "industrial robots, not elsewhere specified or included". In broader trade classifications, this category is embedded within the Standard International Trade Classification (SITC) code 72849, which covers "mechanical appliances with individual functions, n.e.s.".

Our reliance on international trade data means we solely observe direct imports by firms. Consequently, this approach does not capture robots purchased indirectly through domestic distributors, or through national producers. As a result, our indicator of robot adoption serves as a lower bound for actual robot investment. This methodology mirrors the assumption used by Acemoglu et al. (2023) in their study of the Netherlands.

To measure worker exposure to automation, we follow Goos et al. (2014) and utilize the Routine Task Index (RTI) linked to 2-digit International Standard Classification of Occupations (2008) (ISCO-08) codes to capture the extent to which tasks performed by a firm's workforce are routine and codifiable. Additionally, as a robustness check, we follow Montobbio et al. (2024) and associate occupations with a PBOE to robotics metric. This metric is constructed by mapping robot-related patents to O*NET task descriptions, which are then linked to 4-digit ISCO-08 codes in the QP dataset.

Our measure of exposure to the minimum wage builds upon the PRCI index proposed by Alexandre et al. (2022), which measures the potential relative cost increase a firm would incur from a minimum-wage rise – essentially, the proportional increase in the firm's wage bill required to lift workers paid below the new floor up to that floor (aggregating the individual wage gaps); see Section 4.4 for the formal definition and formula.

Additional control variables for firm-level regressions include the firm labour productivity, number of employees, firm age, average worker age, share of workers with higher education.[5]

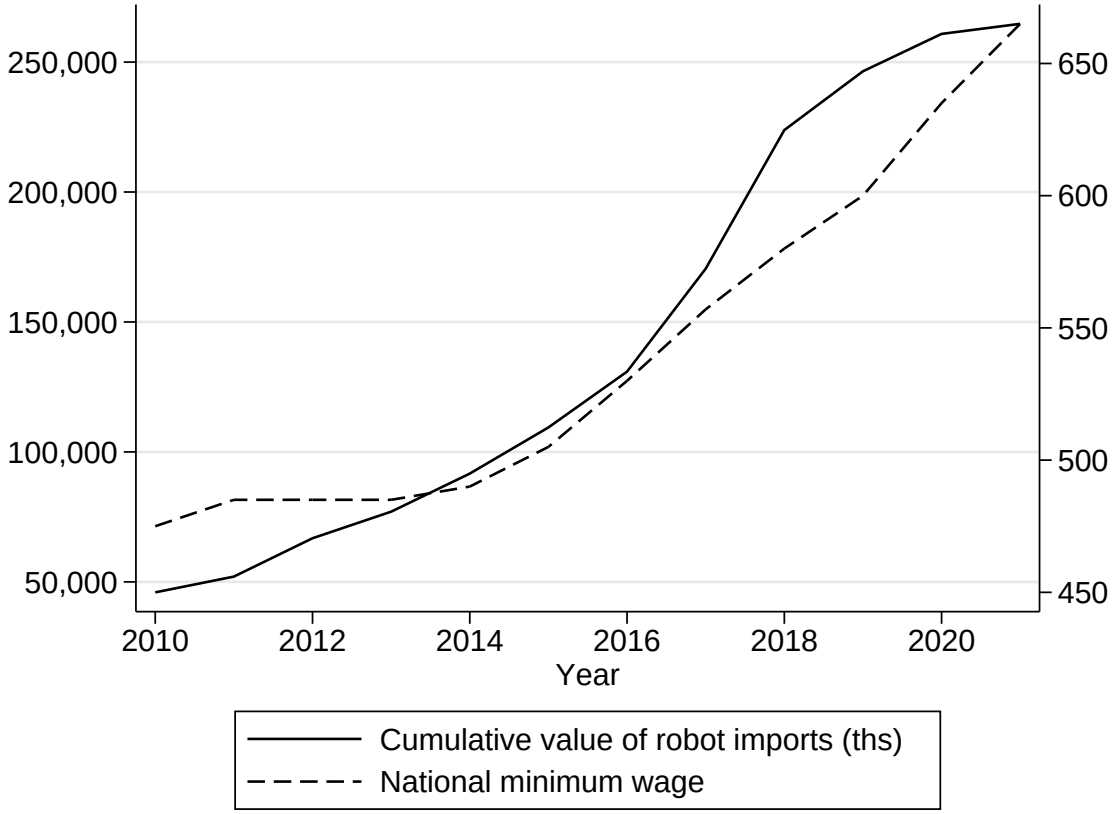
4.3.3 Summary statistics

Although robot-adopting firms constitute a small fraction of the manufacturing firm population, they account for a disproportionately large share of economic activity, representing approximately 19% of total

[4] Industrial robots are programmable manipulators, either fixed or mobile, designed for various industrial automation tasks. These robots typically feature three or more axes and are employed to perform repetitive, hazardous, or precision-demanding tasks. Equipped with advanced technologies such as sensors and cameras, industrial robots can interact with their environment to improve manufacturing quality and productivity. Examples include Cartesian robots for welding and 3D printing, SCARA robots for picking and packing, cyclical or redundant robots for worker assistance, and six-axis robots for material handling.

[5] Robustness checks using the log of the average years of education of the firms' workforce qualitatively confirm our results.

value added and 14% of employment in 2021. Figure 6 shows that robot imports increased steadily over time while the minimum wage rose substantially from 2014 onward.



Notes: *Cumulative robot stock* represents the total value of robot imports in manufacturing sectors, in thousands, adjusted for inflation. *National minimum wage* denotes the monthly (14 months) minimum wage levels in mainland Portugal. *Source:* Authors' calculations based on CI and reports from INE.

Figure 6: Cumulative robot stock in manufacturing and minimum wage levels in Portugal, 2010–2021

Average values of key variables by non-adopters and adopters are presented in Table 22. Further summary statistics by firm size class are available in Table C3 in the Appendix. Two features of this descriptive comparison are important for interpretation. First, robot adopters are strongly positively selected: they are substantially larger, more productive, and more likely to export. Second, consistent with this selection, adopters appear *less* minimum-wage exposed on average (lower PRCI). This does not contradict our Step 1 results (presented in Section 4.4.1) the adoption regressions exploit a predetermined exposure index (fixed at the 2010 workforce/cost structure) and within-firm variation induced by statutory minimum-wage changes, netting out time-invariant firm differences with firm fixed effects. As a result, the causal estimates capture how incremental minimum-wage cost shocks affect adoption within firms, rather than a cross-sectional comparison of high- vs. low-wage firms.

Table C3 in the Appendix reports summary statistics by firm size class, decomposed by adoption status, and reveals substantial heterogeneity across firms. The following patterns stand out. First, larger firms are markedly more productive, employ more educated workers, and are considerably more likely to

Table 21: Variable descriptions and sources

Variable	**Description**	**Source**
Firm-level variables		
Robot adopter	Dummy equal to 1 if the firm has a robot stock greater than zero euros, 0 otherwise	CI
Bought robots	Dummy equal to 1 if the firm purchased robots in t, $t+1$, or $t+2$, 0 otherwise	CI
PRCI	$\text{PRCI}_{it} = \frac{\text{Potential wage bill}_{i,t+1} - \text{Current wage bill}_{it}}{\text{Total costs}_{it}} \times 100$	QP and SCIE
Exporter	Dummy equal to 1 if $\frac{\text{Exports}_{i,t-1}}{\text{Turnover}_{i,t-1}} \times 100 > 5$, 0 otherwise	SCIE
Labour productivity	Value added per worker (in thousands of euros)	SCIE
Number of employees	Total number of workers employed by the firm	SCIE
Tangible assets	Firm's tangible fixed assets (in thousands of euros)	SCIE
Firm age	Years since the firm was established	QP
Average worker age	Average age of the firm's workforce	QP
Share of higher education	Proportion of the firm's workforce with a bachelor's degree or higher	QP
Worker-level variables		
Wages	Logarithm of the worker's real hourly wage	QP
Displacement	Dummy equal to 1 if the worker separates from the firm in year $t+1$, 0 otherwise	QP
Reallocation	Dummy equal to 1 if the worker remains in the same firm but moves to an occupation with strictly lower automation exposure in year $t+1$	QP
Occupation	Worker's 2-digit ISCO-08 occupational category	QP
RTI	Routine Task Index linked to 2-digit ISCO-08 codes, capturing task codifiability and routine intensity (Goos et al. (2014))	QP
PBOE	Patent-Based Occupational Exposure mapping robot-related patents to O^*NET task descriptions (Montobbio et al. (2024))	QP
Log(Educ)	Logarithm of the worker's number of formal years of education	QP

Notes: This table defines the key variables used in the empirical analysis.

Table 22: Summary statistics on adopters and non-adopters

	Non-adopters	Adopters
PRCI	0.24	0.05
RTI	0.20	0.03
PBOE	0.40	0.42
Labour productivity (ths)	27.21	62.61
Number of employees	20.24	211.08
Exporter	0.29	0.79
Tangible assets (ths)	836.26	11259.99
Firm age	17.88	24.90
Worker age	41.56	39.75
Sh. of workers with higher education	0.08	0.19

Notes: The table presents sample averages for firm and worker characteristics, split by robot adoption status. "Non-adopters" refers to firms that never imported industrial robots during the sample period, while "Adopters" refers to firms with a positive robot stock. Variables denoted with "(Thousands (ths))" are expressed in thousands of euros. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

export and to adopt robots, consistent with the presence of fixed adoption costs and scale advantages. In contrast, smaller firms display higher exposure to minimum-wage increases and lower capital intensity, suggesting tighter financial and technological constraints. Second, the positive selection of adopters documented in Table 22 persists within every size class: adopters are consistently more productive and more export-oriented; with the exception of larger firms. The magnitude of this selection gap varies across size classes. Among the smallest firms (fewer than 10 workers), adopters are nearly five times more productive than non-adopters and employ a substantially higher share of university-educated workers, suggesting that small-firm adoption is confined to a narrow set of high-capability units. Among larger firms, the adopter–non-adopter gap narrows considerably, consistent with the idea that scale alone brings firms closer to the adoption margin. Third, minimum-wage exposure (PRCI) declines monotonically with firm size for non-adopters; for adopters, this pattern emerges only among firms with 50 or more employees. Nonetheless, PRCI remains lower for adopters within each size class, reinforcing the view that adoption is driven by firms whose cost structures are less dependent on minimum-wage labour.

4.4 Empirical strategy

This section sets out the paper's two empirical designs. First, we estimate whether predetermined exposure to minimum-wage increases changes firms' probability of purchasing industrial robots. Identification comes from national minimum-wage changes interacted with firms' fixed 2010 wage structures, comparing firms within industry-year cells. Second, we estimate whether robot adoption differentially affects workers according to their pre-adoption task exposure, using worker fixed effects and firm-year fixed effects.

As complementary context, Appendix C.3 documents descriptive firm-level trajectories around adoption (TWFE and event-study DiD) as selection diagnostics only; we do not interpret those patterns causally.

4.4.1 Step 1: Minimum wage and robot adoption

To study whether minimum-wage increases affect robot purchases, we use two complementary approaches. Our baseline specification uses local projections (see Jordà, 2005) with a predetermined exposure measure. This estimator preserves continuous variation in shock intensity and allows the purchase response to unfold over time. As a robustness check, we also estimate a staggered-treatment event study using Callaway and Sant'Anna (2021), mainly to visualise pre-trends and event-time dynamics.

A key challenge in establishing a causal effect of minimum wage increases is in isolating variation in minimum wage exposure that is not driven by the firm's endogenous adjustments. A standard measure of exposure, such as the ratio of minimum wage earners, might change because a firm fires low-wage workers, instead of because the policy changed.

To address this, we construct a *predetermined* measure of exposure. Building on the PRCI index proposed by Alexandre et al. (2022) – itself related to "bite"/exposure measures used in the minimum wage literature (e.g., Draca et al., 2011) – we fix each firm's workforce composition and cost structure at their values in the base year 2010. We then simulate the cost shock imposed by the national minimum wage in year t on this fixed 2010 structure.[6] The predetermined Potential Relative Cost Increase ($PRCI^{pred}$) is defined as:

$$PRCI_{ft}^{pred} = \frac{\sum_{k \in \mathcal{W}_{f,2010}} \max(0, MW_t - w_{k,2010})}{C_{f,2010}} \times 100 \tag{10}$$

where $\mathcal{W}_{f,2010}$ is the set of workers in firm f in 2010, $w_{k,2010}$ is worker k's baseline wage, and $C_{f,2010}$ is the firm's total cost in 2010. To ensure comparability, both the individual worker's baseline wage and the statutory minimum wage (MW_t) are expressed as annual values. Part-time workers are retained in the sample.[7] Variation in this index over time is driven solely by changes in the national minimum wage (MW_t), eliminating simultaneity bias from workforce adjustments.

The key identifying assumption is that, absent minimum-wage changes, firms with different baseline low-wage exposure would have exhibited similar trends in robot purchases within the same industry-year cells. The identifying variation comes from common national minimum-wage changes interacted with predetermined firm-specific exposure based on the 2010 wage structure. Industry-year fixed effects absorb shocks common to firms operating in the same industry and year. As detailed below, the event-study robustness check makes potential differential pre-trends visible.

We estimate the dynamic effect of this shock using local projections. Let $\Delta PRCI_{ft}^{pred} = PRCI_{ft}^{pred} - PRCI_{f,t-1}^{pred}$. For each horizon $h = 0, \ldots, 5$, we estimate:

[6] The base year 2010 was chosen because it is the first year in which the SCIE and QP datasets are harmonised and it precedes the sequence of large minimum wage increases that began in October 2014.

[7] Only 1% of workers in our final sample are part-time workers.

$$Y_{f,t+h} = \beta_h \Delta PRCI_{ft}^{pred} + \gamma_h X_{ft} + \alpha_f + \delta_{st} + \varepsilon_{f,t+h}^{(h)} \tag{11}$$

where $Y_{f,t+h}$ is a dummy variable equal to 1 if firm f purchases robots at any point in the three-year window starting in year $t + h$ (i.e., in years $t + h$, $t + h + 1$, or $t + h + 2$), and 0 otherwise. For $h = 0$, this corresponds to purchases in years t through $t + 2$. We use this forward-looking three-year purchase window to capture lumpy investment and planning lags. The vector of coefficients $\{\beta_h\}$ traces the dynamic response of robot purchases to a minimum-wage-induced cost shock.

In the reported specification, X_{ft} includes time-varying firm characteristics measured in year t: log of average worker age, log labour productivity, and the share of workers with higher education – see Table 21 for details. The model also includes firm fixed effects (α_f) and industry–year fixed effects (δ_{st}). Because some components of X_{ft} may respond contemporaneously to minimum-wage shocks, these controls should be interpreted as conditioning variables in the reported specification rather than as variables required for identification. The identifying variation comes from changes in predetermined exposure within firms, compared across firms in the same industry-year cells. Standard errors are clustered at the firm level. To ensure the robustness of our results, we estimate more defensible variants: a specification excluding X_{ft} entirely, a specification using lagged controls ($X_{f,t-1}$) to mitigate potential bad-control problems, and models employing two–way clustered standard errors at both the firm and year levels.

To further ensure the validity of our empirical design, we implement the following set of robustness tests. First, because the statutory minimum wage rises almost monotonically during 2014 – 2021, the first-difference shock $\Delta PRCI_{ft}^{pred}$ is mechanically related to the level of $PRCI_{ft}^{pred}$. To verify that the results are driven by the incremental shock rather than cumulative exposure, we re-estimate Equation 11 including the lagged level $PRCI_{f,t-1}^{pred}$ alongside $\Delta PRCI_{ft}^{pred}$. Second, we estimate group-time average treatment effects using the estimator proposed by Callaway and Sant'Anna (2021), with not-yet-treated firms as controls and a doubly robust score. Because this estimator requires a binary treatment, we classify a firm as treated in the first year its time-varying predetermined exposure $PRCI_{ft}^{pred}$ crosses a fixed cutoff c, where c is the 75th percentile of the 2010 cross-firm distribution of $PRCI^{pred}$. Firms that never cross the cutoff are never treated, while firms that cross later contribute to the not-yet-treated control group before treatment. We use this discretised design as a robustness check for pre-trends and event-time visualisation rather than as our main estimator. Because $PRCI^{pred}$ is based on the 2010 workforce, this analysis is restricted to firms observed in the base year.

4.4.2 Step 2: Robot adoption and worker-level outcomes

The second step examines how robot adoption affects workers within adopting firms, focusing on heterogeneous effects by workers' pre-adoption exposure to automation. The outcomes are wages, displacement, and occupational reallocation. Identification comes from within-firm, within-year comparisons between workers with different pre-adoption exposure levels.

The key identifying assumption is that, absent robot adoption, workers with different pre-adoption exposure within the same firm would have experienced similar outcome trends. Because firm-year fixed effects absorb any average effect of adoption common to all workers within a firm-year, the design identifies only how adoption reshapes the distribution of outcomes across exposure levels within adopting firms.

Exposure is frozen at a pre-adoption baseline to limit endogeneity from post-adoption reallocation. We utilise the Routine Task Index (RTI), linked to 2-digit ISCO-08 codes, to capture susceptibility to automation driven by task codifiability. Because this classification is relatively coarse, the credibility of this design depends on sufficient within-firm dispersion in exposure. As a robustness check, we also use the PBOE measure proposed by Montobbio et al. (2024), linked to 2-digit ISCO-08 codes. Appendix Table C4 documents that within-firm dispersion in these exposure measures is highly robust. Almost 96% of the manufacturing workers in our 2010 baseline sample were employed in firms with multiple occupational categories, providing ample pre-treatment within-firm variation to identify the interaction effect.

Formally, for worker i employed at firm j in year t, let $Exposure_{ij}^{0}$ denote the worker's baseline exposure. For incumbent workers in adopting firms, this is measured in the last year prior to the firm's robot adoption. If a worker joins an adopting firm after the adoption event has already occurred, their exposure is fixed at their year of entry. For workers in non-adopting firms, exposure is time-variant, as there is no post-adoption reallocation. Freezing exposure means that the interaction between robot adoption and exposure captures differential effects across workers who were initially more or less exposed, rather than differences generated by post-adoption occupation switching. Because exposure remains time-variant for the control group, its standalone main effect is not fully absorbed by worker fixed effects and remains identified in our estimations.

To estimate the effect of robots on workers, we estimate the following high-dimensional fixed effects specification:

$$y_{ijt} = \alpha_i + \phi_{jt} + \beta \left(Robots_{jt} \times Exposure_{ij}^{0} \right) + \gamma X_{it} + \varepsilon_{ijt} \tag{12}$$

where y_{ijt} denotes the outcome of interest (log wages, displacement, or reallocation), α_i are worker fixed effects, and ϕ_{jt} are firm-year fixed effects in the most restrictive specification. $Robots_{jt}$ indicates whether firm j has adopted robots by year t, and $Exposure_{ij}^{0}$ is fixed at the last pre-adoption year. The vector X_{it} includes worker-level characteristics: the log of worker age and the log of years of formal schooling. We report progressively restrictive specifications (Ordinary Least Squares (OLS); worker and firm fixed effects; and finally worker and firm-year fixed effects) for transparency. In the preferred specification with worker and firm-year fixed effects, the parameter of interest is the interaction coefficient β, which captures whether workers with higher pre-adoption exposure experience different post-adoption outcomes than less exposed co-workers within the same firm-year. Because firm-year fixed effects absorb any common adoption effect at the firm-year level, this specification does not separately identify the average effect of robot adoption on all workers. In the current reported results, standard errors are clustered at the firm level.

In the preferred specification, identification comes from comparing workers with different pre-adoption

exposure within the same firm-year. A negative β implies that, within adopting firms, workers initially in more exposed tasks experience relatively worse outcomes than their less exposed colleagues.

For wage outcomes, we estimate this equation on the sample of "firm and occupation stayers", that is, workers who remain employed at the firm before and after the adoption event and in the same occupation, to limit composition changes around adoption. This restriction may attenuate the estimated wage effect if the most negatively affected high-RTI workers are more likely to exit, so wage coefficients should be interpreted together with the separation estimates rather than as average effects for all workers.

For displacement, $Exit_{ij,t+1} = 1$ if worker i is no longer employed by firm j in year $t+1$. We exclude observations for which firm j is not observed in $t+1$. The coefficient β therefore captures whether, within the same firm-year, more exposed workers are more likely to separate following adoption than less exposed workers. We estimate this model separately for fixed-term workers because firing costs differ by contract type.

For reallocation, $Reallocation_{ijt} = 1$ if worker i remains in firm j but moves by $t+1$ to an occupation with strictly lower automation exposure than the previous one. We exclude workers who switch firms between t and $t+1$. This outcome captures downward task reassignment within the firm rather than general mobility. As a robustness check, we also redefine reallocation as within-firm occupation switch, regardless of whether task exposure falls. This helps distinguish "no downward-exposure moves" from "no internal mobility at all".

Connecting Steps 1 and 2. The two steps address different causal questions. Step 1 asks whether minimum-wage-induced cost pressure changes the likelihood of robot purchase. Step 2 asks, conditional on adoption occurring, whether workers with different pre-adoption exposure are affected differently within the same firm-year. We do not estimate the mediated effect of minimum wages on worker outcomes through robots. The credibility of Step 2 rests on exposure being measured before adoption and on the absence of differential within-firm outcome trends by exposure in the absence of adoption.[8]

4.5 Results

This section presents the results and discussion from our empirical strategy. We structure the analysis in two parts. First, in subsection 4.5.1, we examine the dynamic impact of minimum wage shocks on the decision to adopt industrial robots. We find that minimum wage shocks have a significant but delayed effect on adoption. Second, in subsection 4.5.2, we provide causal evidence on how robot adoption affects worker outcomes, including wages, displacement, and reallocation. We find that workers with higher exposure to automation experience relative wage penalties among stayers, and that displacement and limited internal reallocation are concentrated among workers on fixed-term contracts. As complementary context, Appendix C.3 reports descriptive firm-level trajectories around adoption and shows that future

[8] Instrumenting Step 2 with Step 1 exposure is unlikely to be informative because robot adoption is rare in the data, implying a weak first stage.

adopters are positively selected; we use this evidence as a selection diagnostic rather than a basis for causal interpretation of firm performance.

4.5.1 Step 1: Minimum wage and robot adoption

The results from Equation 11 are reported in Figure 7 and Table 23. Our results show that minimum wage shocks have a delayed response effect on the adoption of industrial robots, with a positive effect only being observed at $t + 3$, and continuing to increase over time. In $t + 5$, a one-percentage-point increase in the predetermined minimum-wage cost shock $\Delta PRCI_{ft}^{pred}$, measured in percentage points of 2010 total cost, is associated with a 0.19 percentage point increase in the likelihood of adopting robots within a three-year window.

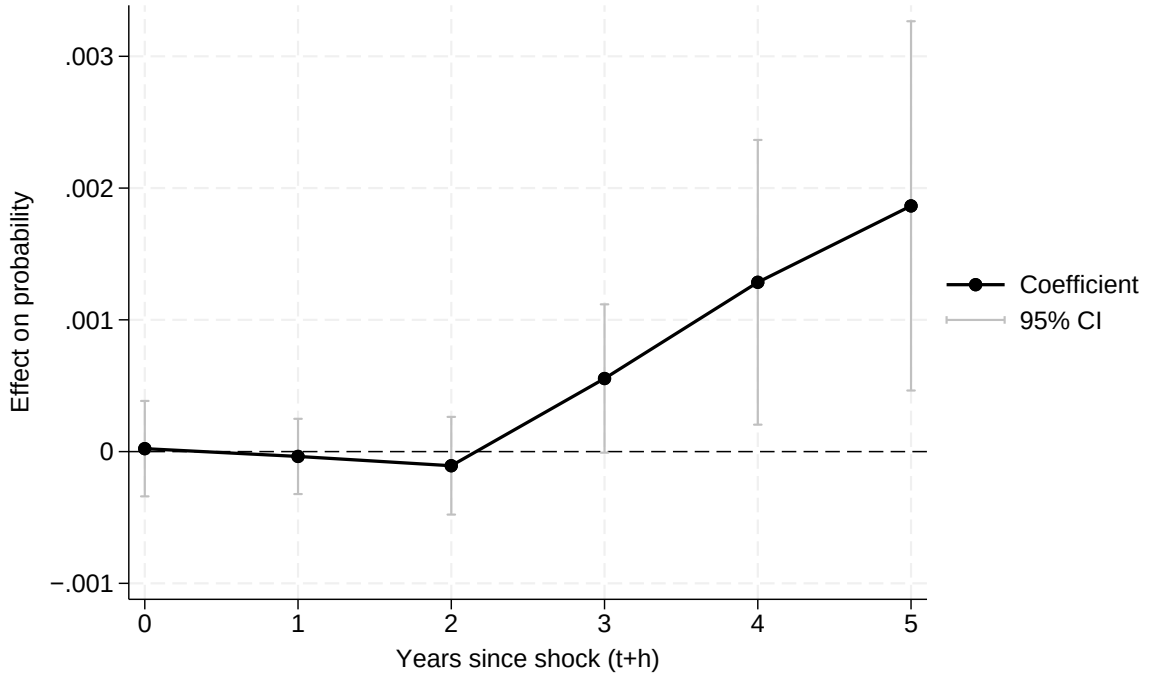
To understand magnitudes, we compare the estimated effect to both the dispersion of the shock and the baseline likelihood of adoption. The standard deviation of the cost shock $\Delta PRCI^{pred}$ is 0.5047, while the average probability of purchasing robots in any given three-year window is 0.51% (0.0051). Using the $t+5$ coefficient from Table 23 (0.0019), a one standard deviation shock increases the adoption probability by $0.0019 \times 0.5047 = 0.00096$, i.e. about 0.096 percentage points. In relative terms, this corresponds to an 18.8% increase compared to the baseline mean ($0.00096/0.0051$). Thus, while robot adoption is a rare event, the estimated effect is economically meaningful relative to the baseline adoption probability.

To provide context on the effective sample size driving these estimates, we observe that 235 firms (out of 377 robot adopting firms in 2021) switch from being non-adopters to adopters during the 2010 to 2021 period. Notably, 123 of these switching firms are located in the 75th percentile of the 2010 cross-firm distribution of $PRCI^{pred}$. This concentration highlights that our identifying variation is driven by a meaningful number of transitioning firms, the majority of which were highly exposed to the minimum wage shocks.

We perform several robustness checks to validate these findings. As the number of observations drops as the horizon h increases, we report results with a stable sample in Table C5. Furthermore, Figure C1 reports results using a stricter definition of the dependent variable, considering robot purchases only in year t rather than in the three-year window. These alternative specifications qualitatively confirm the baseline results.

Table C5 suggests a non-monotonic dynamic pattern. In the stable sample, the coefficients are negative and statistically significant at t and $t+1$, close to zero at $t+2$, positive but imprecisely estimated at $t+3$, and positive and statistically significant at $t+4$ and $t+5$. This pattern is consistent with an initial tightening of investment capacity followed by medium-run substitution towards automation, but it should be interpreted as suggestive robustness evidence rather than definitive proof of that mechanism. This non-monotonicity nevertheless differentiates our findings from Fan et al. (2021), who document gradual but monotonically increasing adoption responses in the Chinese context, and is consistent with models in which firms face binding financial constraints that interact with investment incentives.

To further test the validity of our identification strategy, Figure 8 presents event-study estimates using



Notes: The figure plots the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t+h$ to $t+h+2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). See section 4.4 for further details on the specification. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. See Table 23 for further details. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure 7: Minimum wage and robot adoption – local projection coefficients

the estimator proposed by Callaway and Sant'Anna (2021). The results confirm our baseline findings. First, we find no evidence of differential pre-trends (coefficients $t-5$ to $t-1$ are statistically indistinguishable from zero), supporting the parallel trends assumption. Second, the treatment effect on robot adoption is positive and statistically significant (in $t+3$ and $t+4$ at a 5% statistical significance level), but materialises with a lag. We observe a significant increase in the probability of adoption starting 3 to 4 years after the firm first crosses the fixed high-exposure cutoff.

Additional robustness checks for Equation 11 considering the lagged level $PRCI_{f,t-1}^{pred}$ alongside $\Delta PRCI_{ft}^{pred}$; employing two–way clustered standard errors at both the firm and year levels; excluding X_{ft} entirely; and a specification using lagged controls ($X_{f,t-1}$) are present in the Appendix Tables C6, C7, C8, and C9.

Before turning to worker-level outcomes, it is worth emphasising that robot adoption is an endogenous strategic choice. Appendix C.3 documents descriptive firm-level trajectories around adoption. The event-study patterns reveal that future adopters already exhibit stronger pre-adoption trajectories, indicating positive selection into automation; accordingly, we do not interpret firm-level correlations as causal. This motivates our focus on worker outcomes identified from within-firm comparisons in Step 2.

Table 23: Minimum wage and robot adoption

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI_{it}^{pred}$	0.0000	-0.0000	-0.0001	0.0006*	0.0013**	0.0019***
	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0006)	(0.0007)
Log(Worker age)	-0.0040**	-0.0013	0.0018	0.0023	0.0004	0.0027
	(0.0020)	(0.0022)	(0.0023)	(0.0026)	(0.0025)	(0.0028)
Log(productivity)	0.0005	0.0000	-0.0005	0.0003	0.0000	0.0001
	(0.0004)	(0.0005)	(0.0006)	(0.0006)	(0.0006)	(0.0007)
% bachelor	0.0031	0.0035	0.0044	0.0065*	0.0078**	0.0005
	(0.0029)	(0.0034)	(0.0036)	(0.0037)	(0.0037)	(0.0036)
Observations	172092	150317	132200	114930	98434	82565

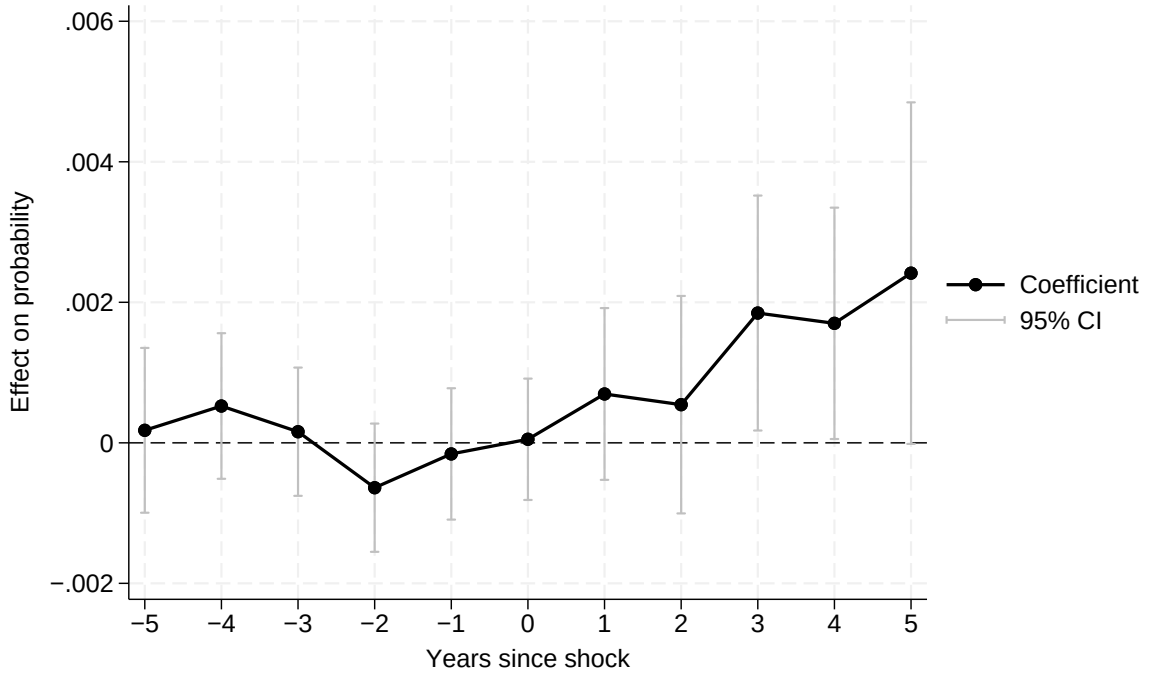
Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t+h$ to $t+h+2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). All specifications include firm fixed effects, industry-year fixed effects, and time-varying firm controls (log of average worker age, log labour productivity, and the share of workers with higher education). Standard errors clustered at the firm level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

4.5.2 Step 2: Worker level outcomes

Results from Equation (12), with log hourly wages as the outcome, are presented in Table 24. These estimates are computed on the sample of firm-and-occupation stayers (see Section 4.4), so they capture differential wage changes among workers who remain employed in the same firm and occupation around adoption. Because the wage regression is estimated on firm-and-occupation stayers with worker fixed effects and firm-year fixed effects, the interaction coefficient should be interpreted as a relative wage effect among retained co-workers within the same firm-year, not as the average wage effect of robot adoption on all workers.

In our preferred specification (column 4), the interaction term $Robots \times RTI$ is negative and statistically significant. A one standard deviation increase in RTI (approximately 0.5) implies an additional relative wage penalty of about 0.6% (0.5×0.0111) among stayers within adopting firms. To contextualise this magnitude, the within-firm standard deviation of log wages across workers in our sample is 0.35. Therefore, the automation-induced penalty represents a highly targeted, modest compression of approximately 1.6% of a standard deviation in within-firm wage dispersion.[9] Furthermore, a descriptive calculation shows that workers in the bottom quartile of the baseline wage distribution experienced average real wage growth of approximately 40 log points between 2010 and 2021. Consequently, while the relative penalty for high-RTI tasks is statistically robust, it is quantitatively small when compared to the direct, first-order wage gains driven by the minimum wage trajectory.

[9] $0.5 \times 0.0111 / 0.3517 = 0.01578$ (or 1.6%)



Notes: The figure plots event-study coefficients estimating the dynamic effect of minimum wage exposure on robot adoption using the Callaway & Sant'Anna (2021) estimator. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in a three-year window. Event time zero ($g = 0$) is defined as the first year a firm's predetermined exposure ($PRCI_{ft}^{pred}$) crosses the 75th percentile of the baseline 2010 cross-firm distribution. The control group consists of not-yet-treated firms. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure 8: Minimum wage and robot adoption – differences in differences estimation

Columns (1) – (3) also reveal important insights regarding selection and sorting. Column (1) indicates that, on average, workers in robot adopting firms earn a wage premium of about 4% (not statistically significant). However, the results in columns (2) and (3) suggests that this premium is driven by unobserved worker and firm heterogeneity, that is, "better" firms buy robots and hire "better" workers, as shown by the decrease in the coefficient associated with "Robots" in columns (2) and (3). Furthermore, we find that workers performing routine tasks are, on average, worse paid, regardless of whether their firm adopts robots (columns 1–4). Finally, as already noted from the results in column (4), consistent with the substitution hypothesis, this routine-wage penalty is aggravated upon adoption. This suggests that, within the sample of firm-and-occupation stayers, more exposed workers experience less favourable wage growth than less exposed co-workers within the same firm-year after adoption.

Figure 9 plots the dynamic evolution of relative wages for more exposed versus less exposed firm-and-occupation stayers around robot adoption.[10] The estimated coefficients indicate no systematic differential pre-trends in the periods leading up to adoption; specifically, coefficients for $t-4$ and $t-3$ are statistically indistinguishable from zero. We observe one marginally significant negative coefficient at $t-2$; taken in isolation this does not provide strong evidence against parallel trends, but it may be consistent with early

[10] Exposure is proxied by the RTI.

Table 24: Robot adoption, exposure, and worker wages

	(1)	(2)	(3)	(4)
Robots=1	0.0417	0.0087	0.0087	
	(0.0285)	(0.0064)	(0.0064)	
RTI	-0.1184***	-0.0152***	-0.0152***	-0.0153***
	(0.0038)	(0.0014)	(0.0014)	(0.0014)
Robots=1 × RTI	-0.0949***	-0.0114**	-0.0114**	-0.0111**
	(0.0197)	(0.0054)	(0.0054)	(0.0054)
Log(Worker age)	0.6149***	1.0568***	1.0568***	0.9887***
	(0.0103)	(0.0242)	(0.0242)	(0.0270)
Log(Educ)	0.5042***	0.0110***	0.0110***	0.0116***
	(0.0086)	(0.0020)	(0.0020)	(0.0016)
Observations	3164488	3151859	3151859	3099607
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is the log of hourly wages. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

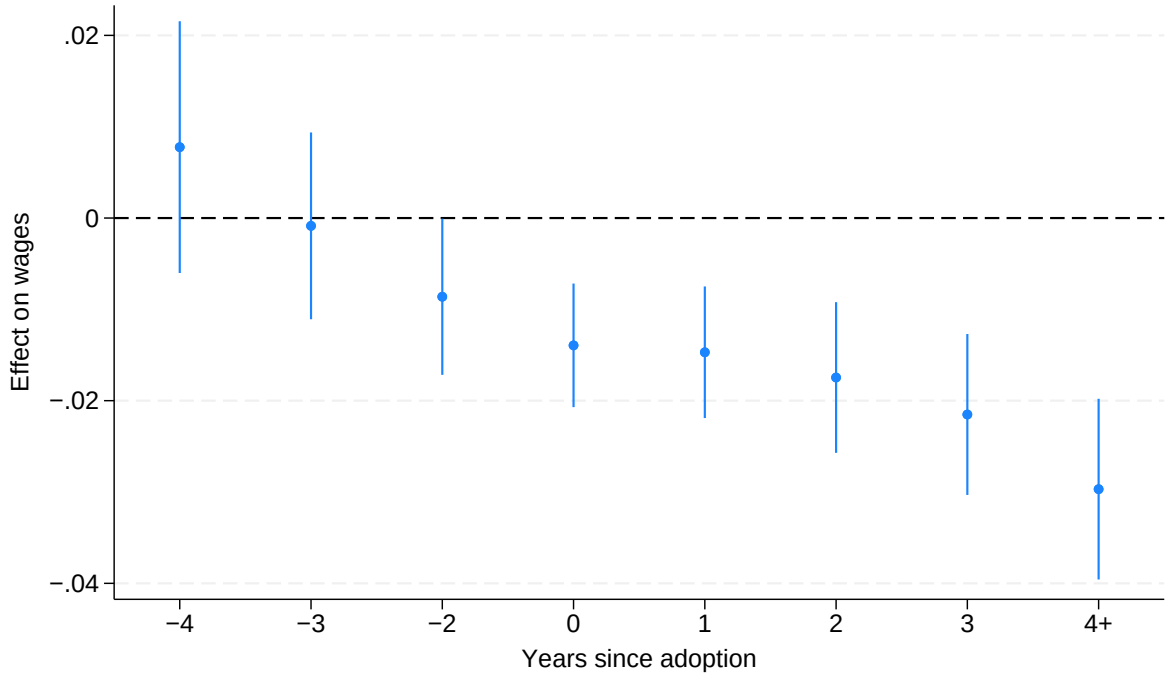
restructuring in anticipation of the investment decision.

Importantly, a clear structural break occurs at the time of adoption ($t = 0$). The relative wage penalty among stayers becomes statistically significant immediately upon adoption and grows over time, consistent with slower wage growth for more exposed retained workers within adopting firms. Because workers who separate from the firm or switch occupation are excluded from this sample, these estimates should be read jointly with the separation results below.

Given their lower firing costs, fixed-term workers are especially vulnerable when firms adjust employment following automation. For this reason, we estimate Equation (12) for a subsample of fixed-term workers in Table 25, for the entire sample in Table 26, and (to align the evidence with the dual labour market mechanism) for open-ended contract workers in Table 27.

The results from our preferred specification (columns 4) show that workers with high exposure to automation (high RTI) are significantly more likely to be displaced from robot adopting firms, but this effect is concentrated among workers with fixed-term contracts. A one standard deviation increase in the RTI is associated with a 1.30 p.p. (0.5×0.0259) increase in the likelihood of displacement for these workers. To gauge economic significance, the baseline annual exit probability for fixed-term workers in the estimation sample is 6.4%. Hence, the 1.30 p.p. effect corresponds to an increase of approximately 20.4% relative to baseline.

Furthermore, columns (1)–(3) suggest that while workers in robot adopting firms are, on average, less



Notes: The figure plots the event-study coefficients representing the dynamic wage effects of robot adoption, differentiated by workers' pre-adoption automation exposure. The dependent variable is the log of hourly wages. Event time zero ($t = 0$) is defined as the year of the firm's first observed robot purchase. See section 4.4 for further details on specifications and sample. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure 9: Robot adoption, exposure, and worker wages – event study

likely to be displaced (Table 25 and 26), when considering all workers, those performing routine tasks face a higher baseline risk of displacement (columns 2 – 4 of Table 26). [11] These results suggest that some degree of substitution of human capital by physical capital occurs when firms automate, with the burden of adjustment falling disproportionately on fixed-term contract workers.

Figure 10 presents the event-study estimates for worker displacement, for fixed-term workers. The coefficients for the pre-adoption periods ($t-4$ to $t-2$) are statistically indistinguishable from zero, providing strong support for the parallel trends assumption. In contrast, we observe a sharp structural break at the time of adoption ($t = 0$). The probability of displacement for high-exposure workers rises significantly immediately upon the introduction of robots and remains elevated throughout the post-adoption period.

Results from Equation 12, considering worker reallocation within the same firm as the outcome, are presented in Tables 28 and 29 for fixed-term and all workers, respectively, and in Table 30 for open-ended contract workers.

Recall that reallocation is defined as a within-firm occupation switch to a role with strictly lower automation exposure (a "downward exposure" move). In our preferred specification (column 4), fixed-term

[11] Note that when considering all workers, the baseline separation rate lowers from 6.4% as seen for fixed-term workers to 3.1%

Table 25: Robot adoption, exposure, and worker displacement – fixed term workers

	(1)	(2)	(3)	(4)
Robots=1	-0.0098**	-0.0109	-0.0162**	
	(0.0040)	(0.0072)	(0.0064)	
RTI	-0.0004	0.0032	0.0013	0.0035
	(0.0012)	(0.0034)	(0.0032)	(0.0033)
Robots=1 × RTI	-0.0176***	0.0150	0.0197**	0.0259**
	(0.0038)	(0.0102)	(0.0096)	(0.0104)
Log(Worker age)	-0.0257***	0.1225***	0.2411***	0.2366***
	(0.0022)	(0.0269)	(0.0296)	(0.0295)
Log(Educ)	0.0089***	-0.0213***	-0.0095	-0.0105
	(0.0017)	(0.0066)	(0.0062)	(0.0064)
Observations	705906	601092	599510	560714
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Only workers under a fixed-term contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

contract workers in adopting firms with high exposure to automation are significantly less likely to be reallocated internally (Table 28). A one standard deviation increase in the RTI (0.5) reduces the probability of protective internal reassignment by approximately 3.02 percentage points (0.5×0.0604). Given the low baseline reallocation rate of 2.17% in the sample, this negative shock is economically massive, essentially eliminating the prospect of downward-exposure mobility for these workers. The coefficients for all workers and open-ended contracts (Tables 29 and 30) are not statistically significant.

The reallocation results should nevertheless be interpreted cautiously. The negative coefficient on $Robots \times RTI$ in Table 28 does not, by itself, identify a firm decision not to retrain. If high-RTI fixed-term workers leave the firm through displacement (Table 25) before any within-firm move can occur, the reallocation sample will mechanically under-represent the most exposed workers. The estimates therefore document limited *observed* internal reassignment for exposed temporary workers after adoption, but they do not separately identify whether firms chose not to retrain them or whether displacement pre-empted reallocation.

Figure 11 plots the dynamic evolution of the reallocation interaction for all workers. The coefficients are close to zero throughout the event window, indicating no differential shift in internal mobility between more and less exposed workers around adoption. In Appendix Figure C7, we report the corresponding event study for fixed-term workers. The coefficients are likewise imprecisely estimated and do not trace a

Table 26: Robot adoption, exposure, and worker displacement – all workers

	(1)	(2)	(3)	(4)
Robots=1	-0.0035**	-0.0056**	-0.0142***	
	(0.0018)	(0.0028)	(0.0018)	
RTI	-0.0028***	0.0039***	0.0025*	0.0032***
	(0.0003)	(0.0012)	(0.0013)	(0.0008)
Robots=1 × RTI	-0.0051***	0.0033	0.0042**	0.0030
	(0.0013)	(0.0026)	(0.0019)	(0.0019)
Log(Worker age)	-0.0399***	0.0752***	0.1932***	0.1640***
	(0.0011)	(0.0087)	(0.0095)	(0.0085)
Log(Educ)	0.0058***	-0.0205***	-0.0076***	-0.0091***
	(0.0006)	(0.0022)	(0.0019)	(0.0017)
Observations	3333196	3221097	3219617	3176998
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

clear dynamic pattern. However, the event-study design is poorly suited to detect the effect identified in Table 28: because displacement removes the most exposed fixed-term workers from the sample before within-firm moves can be observed (Table 25 and Figure 10), the reallocation sample in each post-adoption bin is progressively selected against exactly the workers for whom the static specification detects a negative effect. The event study is therefore uninformative about this margin, and we rely on the pooled interaction in Table 28 for inference. That specification shows that, within adopting firms, fixed-term workers in more automatable tasks are significantly less likely to be reallocated to lower-exposure occupations than their less exposed co-workers in the same firm-year – pointing to a relative disadvantage in protective reassignment for the most exposed temporary workers.

Taken together, the results from Tables 25, 26, 28, and 29 point to a dualised pattern of workforce adjustment. Following robot adoption, high-exposure workers on fixed-term contracts are more likely to separate and less likely to be reallocated to lower-exposure occupations than comparable co-workers within the same firm-year. This joint pattern is consistent with adjustment being concentrated on the temporary workforce. However, because displacement and reallocation draw on overlapping populations – workers who exit cannot subsequently be reallocated – the reallocation deficit may partly reflect survivorship selection rather than a firm decision not to retrain. A plausible interpretation is that stronger protection on open-ended contracts limits comparable adjustment among insiders, but the current results do not

Table 27: Robot adoption, exposure, and worker displacement – open-ended contract workers

	(1)	(2)	(3)	(4)
Robots=1	-0.0009	-0.0060**	-0.0134***	
	(0.0015)	(0.0027)	(0.0019)	
RTI	-0.0035***	0.0040***	0.0019	0.0021***
	(0.0003)	(0.0012)	(0.0013)	(0.0007)
Robots=1 × RTI	-0.0022**	0.0010	0.0001	-0.0009
	(0.0010)	(0.0024)	(0.0018)	(0.0016)
Log(Worker age)	-0.0282***	0.1454***	0.2484***	0.2032***
	(0.0009)	(0.0101)	(0.0108)	(0.0095)
Log(Educ)	0.0051***	-0.0168***	-0.0041**	-0.0046***
	(0.0005)	(0.0023)	(0.0020)	(0.0016)
Observations	2627290	2544044	2542961	2497318
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Only workers under an open-ended contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

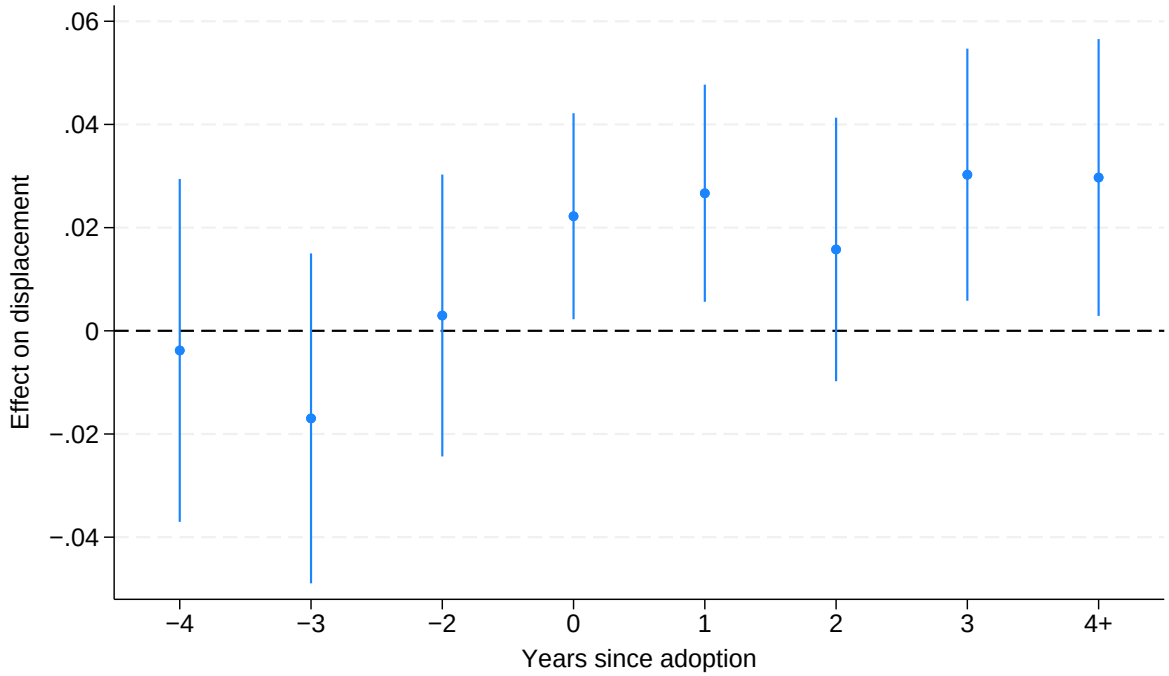
Source: Authors' calculations based on CI, QP, and SCIE.

directly identify that margin.

All interaction coefficients from our preferred specifications are robust to the inclusion of two-digit occupation fixed-effects, in coefficient size, sign and statistical significance. However, the inclusion of this parameter is highly collinear with our exposure measure, with identification being driven by workers in never adopting firms that switch occupation. For this reason, we do not report this additional specification.

In Tables C11, C12, C13, C14, and C15 we report results using the alternative exposure (PBOE) measure proposed by Montobbio et al. (2024), which maps robot-related patents to occupational task descriptions. These results qualitatively confirm our baseline findings: the wage penalty for more exposed workers, the concentration of displacement among fixed-term workers, and the limited internal reallocation are all robust to this alternative operationalisation of task exposure. When comparing the RTI and patent-based exposure results, the raw interaction coefficients should not be compared directly because the two exposure measures are scaled differently. Any quantitative comparison should therefore be based on standardised contrasts, such as a one-standard-deviation increase or a p75-p25 difference. At this stage, the alternative exposure measure is best interpreted as a qualitative robustness check on sign, significance, and the overall pattern of heterogeneous effects.

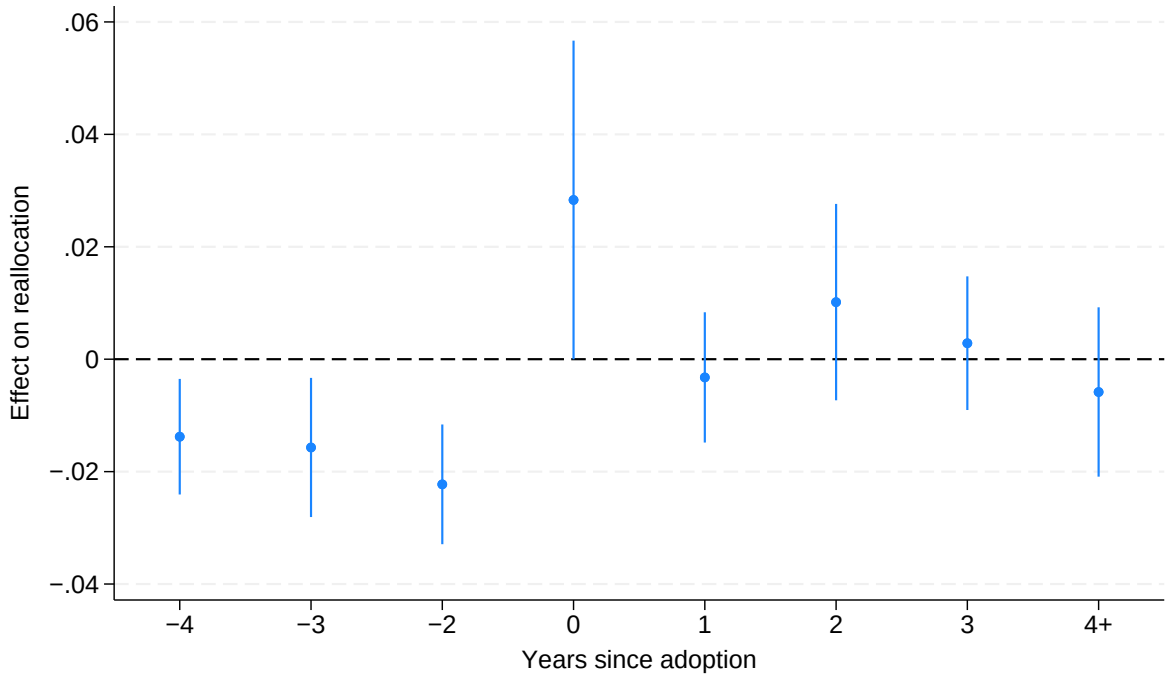
Finally, Tables C16, C17, and C18 report the effects of robot adoption on general occupational reallocation, defined as a switch to any other occupation within the firm, irrespective of the new role's



Notes: The figure plots the event-study coefficients representing the dynamic displacement effects of robot adoption, differentiated by workers' pre-adoption automation exposure. The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Event time zero ($t = 0$) is defined as the year of the firm's first observed robot purchase. See section 4.4 for further details on specifications and sample. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level.
Source: Authors' calculations based on CI, QP, and SCIE (INE).

Figure 10: Robot adoption, exposure, and worker displacement – fixed term workers, event study

automation exposure. For all three subsamples (fixed-term, all, and open-ended contract workers), the interaction coefficient is statistically insignificant. This null result helps sharpen the interpretation of Table 28: the reduction in internal mobility for exposed fixed-term workers is specific to downward-exposure moves rather than reflecting a general freeze on all internal occupational switches. In other words, it is not the case that adopting firms immobilise their temporary workforce altogether; rather, they do not reassign exposed fixed-term workers towards less automatable tasks. Read jointly with the displacement evidence, this suggests that the relevant adjustment margin for these workers is exit from the firm, not protective reallocation within it.



Notes: The figure plots the event-study coefficients representing the dynamic displacement effects of robot adoption, differentiated by workers' pre-adoption automation exposure. The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Event time zero ($t = 0$) is defined as the year of the firm's first observed robot purchase. See section 4.4 for further details on specifications and sample. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure 11: Robot adoption, exposure, and worker reallocation – all workers, event study

Table 28: Robot adoption, exposure, and worker reallocation – fixed term

	(1)	(2)	(3)	(4)
Robots=1	0.0031	0.0092**	0.0046	
	(0.0038)	(0.0039)	(0.0058)	
RTI	0.0174***	0.2288***	0.2865***	0.2741***
	(0.0010)	(0.0075)	(0.0090)	(0.0093)
Robots=1 × RTI	0.0073*	-0.0496***	-0.0512***	-0.0604***
	(0.0043)	(0.0074)	(0.0085)	(0.0102)
Log(Worker age)	-0.0034**	-0.0132	-0.0154	0.0054
	(0.0014)	(0.0165)	(0.0180)	(0.0187)
Log(Educ)	0.0120***	0.0059	0.0050	0.0021
	(0.0009)	(0.0038)	(0.0046)	(0.0045)
Observations	659942	556927	555822	518391
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Only workers under a fixed-term contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table 29: Robot adoption, exposure, and worker reallocation – all workers

	(1)	(2)	(3)	(4)
Robots=1	0.0033	0.0005	-0.0011	
	(0.0027)	(0.0030)	(0.0035)	
RTI	0.0104***	0.1665***	0.1802***	0.1737***
	(0.0005)	(0.0037)	(0.0039)	(0.0040)
Robots=1 × RTI	0.0023	-0.0042	-0.0026	-0.0027
	(0.0026)	(0.0037)	(0.0037)	(0.0036)
Log(Worker age)	-0.0080***	0.0005	-0.0069	0.0005
	(0.0010)	(0.0106)	(0.0110)	(0.0086)
Log(Educ)	0.0079***	0.0045**	0.0028	0.0008
	(0.0008)	(0.0018)	(0.0020)	(0.0014)
Observations	3231694	3120141	3119025	3077724
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table 30: Robot adoption, exposure, and worker reallocation – open-ended contract workers

	(1)	(2)	(3)	(4)
Robots=1	0.0036	-0.0004	-0.0015	
	(0.0027)	(0.0035)	(0.0038)	
RTI	0.0091***	0.1725***	0.1837***	0.1787***
	(0.0005)	(0.0039)	(0.0042)	(0.0047)
Robots=1 × RTI	0.0009	0.0023	0.0044	0.0054
	(0.0023)	(0.0039)	(0.0040)	(0.0038)
Log(Worker age)	-0.0070***	0.0134	0.0074	0.0136
	(0.0009)	(0.0109)	(0.0114)	(0.0086)
Log(Educ)	0.0071***	0.0030	0.0018	0.0000
	(0.0009)	(0.0021)	(0.0023)	(0.0017)
Observations	2571752	2489841	2489003	2444735
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Only workers under an open-ended contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

4.6 Conclusion

This paper studies two related questions: whether minimum wage increases induce firms to adopt industrial robots, and how robot adoption reshapes worker outcomes within firms in a segmented labour market. Using linked employer–employee data and information on industrial robot imports for Portuguese manufacturing firms between 2010 and 2021, we estimate the dynamic effect of predetermined exposure to statutory minimum wage increases on robot adoption and, separately, heterogeneous worker-level effects of adoption within firms. As complementary context, we also document descriptive firm-level trajectories around adoption.

Our first main result is that minimum-wage-induced labour cost shocks increase the probability of robot adoption, but only with a substantial delay. The response emerges only after roughly three to five years, consistent with planning, financing, and organisational frictions in automation investment. This implies that assessments of minimum wage policy that focus only on contemporaneous employment effects may miss an important medium-run technology margin.

Our second main result is that automation reshapes workers' outcomes unevenly within adopting firms. Among workers who remain in the same firm and occupation, those in more automatable tasks experience less favourable wage growth than less exposed co-workers within the same firm-year. At the same time, job separations rise disproportionately among high-exposure workers on fixed-term contracts, and observed within-firm reallocation for that group is limited. These wage estimates are therefore relative effects among stayers, not average wage effects for all workers. The joint pattern points to a dualised form of adjustment concentrated on temporary workers. It is consistent with stronger protection for insiders, but the current evidence does not directly identify the open-ended contract margin, nor does the reallocation evidence separately distinguish a firm decision not to retrain from displacement that pre-empts reassignment.

Appendix C.3 helps clarify how these results should be read. Firms that go on to adopt robots already display stronger pre-adoption trajectories in productivity, scale, and exporting, so correlations between adoption and firm performance are treated as descriptive selection diagnostics rather than as causal effects of automation on firm outcomes. This is precisely why the worker analysis relies on within-firm comparisons across workers with different pre-adoption exposure, rather than on firm-level before-and-after comparisons.

Taken together, the findings show that labour-cost shocks can accelerate automation, but that this adjustment is gradual and its incidence within the firm is uneven. The paper therefore links a policy-driven incentive to automate to the distribution of adjustment costs inside firms, without claiming to estimate the mediated effect of minimum wages on worker outcomes through robots. In segmented labour markets, policies that affect labour costs may interact with technology adoption and employment protection in ways that shift adjustment onto more precarious workers.

From a policy perspective, the results suggest that the evaluation of minimum wage increases should account for medium-run technological adjustment as well as for its distribution across workers. In settings

characterised by strong employment protection and labour market segmentation, measures that facilitate worker adjustment – such as training, reskilling, and support for transitions out of highly automatable tasks – may shape who bears the costs of automation. Future research could extend this analysis by examining other forms of automation, exploring general equilibrium effects across sectors, or studying how complementary policies interact with firms' technology choices to shape the distribution of gains from technological change.

5

Conclusion

This thesis has traced firm adjustment along three margins – the timing of adoption of digital technologies, the recomposition of the workforce following a labour-supply shock, and the adoption of industrial robots following a labour-cost shock – using administrative firm-level accounts, linked employer–employee records, and international trade data for Portugal. This chapter summarises what the three empirical chapters jointly establish, states what they deliberately do not claim, restates the thesis-level contributions, draws chapter-specific policy implications, and discusses limitations and avenues for further research.

Chapter 2 establishes that the returns to digital adoption depend on the complexity of the technology and on the human capital that firms can mobilise. Firms with more educated managers and workers, and shorter managerial tenure, are more likely to adopt cloud computing and big data, with the managerial-education effect more pronounced for big data. The performance implications differ across technologies. For cloud computing, early adoption is associated with gains in value added, turnover, employment and productivity relative to non-adopters, but these gains do not differ significantly between first movers and followers, and complementarities with human capital are not robust in explaining firm performance. For big data, early adoption yields productivity and value-added gains that do not dissipate over the sample period, and these gains are amplified by higher managerial and workforce education. The chapter is therefore consistent with a reading in which a standardised technology generates transient first-mover rents, while a capability-intensive technology generates persistent advantages conditional on complementary human capital.

Chapter 3 documents a strong response to the 2009 compulsory schooling reform in the employment of minors in manufacturing firms: in more exposed firms, their share falls by 2.2–3.2 percentage points and the probability of employing any minor declines by 24–29 percentage points. The chapter also documents a differential post-reform pattern of firm adjustment in labour force composition and robot imports. The clearest positive baseline coefficient appears on complementary ICT employment (1.2–1.6 percentage points higher in more exposed firms); an occupation-based proxy is also positive but less precisely estimated; firm-level robot-import indicators remain close to zero. The chapter is explicit that the ICT result is not reform-specific in a clean sense: age-based exposure placebos produce ICT coefficients larger than the baseline, rather than smaller; fake-timing exercises yield small positive pseudo-effects in

some pre-reform years; and trend-augmented and matched specifications attenuate the ICT coefficient sharply. Whether the ICT pattern reflects the reform itself or a broader process of workforce upgrading among firms that relied more on young labour therefore remains an open question. The chapter provides evidence on differential post-reform firm adjustment rather than on schooling-induced automation.

Chapter 4 establishes that firms more exposed to statutory minimum wage increases are more likely to import industrial robots, but that the response materialises only with a delay of three to five years, consistent with lumpy investment, planning and financing frictions (in particular investment costs). Within firms that adopt, the chapter identifies a distributional pattern: among workers who remain in the same firm and occupation, those in more automatable tasks experience less favourable wage growth than less exposed co-workers within the same firm-year. These estimates are relative effects among stayers, not average wage effects for the workforce. In parallel, job separations rise disproportionately among workers on fixed-term contracts with high task exposure, while observed internal reallocation for that group is limited; permanent workers adjust more often through internal reallocation. Firms that adopt robots already display stronger pre-adoption trajectories in productivity, scale and exporting, and the chapter therefore treats the firm-level correlations around adoption as selection diagnostics that motivate the within-firm worker strategy, rather than as causal effects of robots on firm performance.

Three cross-cutting findings emerge when the chapters are read together. First, capabilities and technology complexity jointly determine the returns to adoption: the returns to big data depend on complementary human capital, while the returns to cloud computing do not. Second, firms' responses to labour market pressures are not uniform. A narrow supply-side restriction (the 2009 reform) operates primarily through workforce recomposition, not through capital adoption; a broader cost-side shock (the minimum wage increases of 2014–2021) is associated with robot adoption at the extensive margin, but only after a three- to five-year lag. The two patterns are consistent when viewed jointly: with a small affected input and fixed costs of capital adoption, workforce reshuffling dominates; with a pervasive cost shock, robot adoption becomes worthwhile, but frictions delay it. Third, the incidence of automation within adopting firms is uneven: workers in more automatable tasks face relatively slower wage growth among stayers, and the probability of separation rises disproportionately for fixed-term workers. The incidence reflects task exposure but also the segmented structure of the Portuguese labour market, which shifts the adjustment margin onto the less protected contractual segment.

Three explicit non-claims follow from the evidence. First, Chapter 3 does not provide clean causal evidence that the 2009 compulsory schooling reform induced automation or ICT adoption; the placebo diagnostics are inconsistent with a reform-only reading of the ICT coefficient. Second, Chapter 4 does not estimate the mediated effect of minimum wage increases on worker outcomes through robots; it links a policy-driven incentive to automate (first part) to the within-firm distribution of adjustment costs (second part), but does not combine them into a single mediation estimate. Third, the firm-level trajectories around robot adoption in Chapter 4 are not interpreted as causal effects of robots on firm outcomes: they reflect positive selection into adoption and serve as selection diagnostics that motivate the within-firm worker analysis.

The thesis makes three contributions. Empirically, it links four Portuguese administrative datasets – SCIE, QP, IUTICE and CI – through a common firm identifier to study adoption decisions, responses to a labour-supply shock and responses to a labour-cost shock within a single country and using a common data spine. Methodologically, it combines firm fixed-effects regressions with propensity-score reweighting in Chapter 2, applies the Callaway and Sant'Anna (2021) continuous-treatment Difference-in-Differences estimator to minimum-wage-induced exposure in Chapter 4, and subjects the common-timing Difference-in-Differences design of Chapter 3 to a battery of placebo exercises – age-based exposure placebos, fake-timing, trend-augmented specifications and matched designs – that discipline the interpretation of the baseline estimates and produce an honest characterisation of their limits. Conceptually, it identifies three mediators – technology complexity, the scale of the labour-market shock, and contractual segmentation – through which firms translate labour market pressures into differentiated adjustments, and establishes that automation is not a uniform response to such pressures.

The policy implications of the thesis are threefold. First, Chapter 2 suggests that the effectiveness of policies promoting digital adoption depends on firms' absorptive capacity: public support for advanced technologies such as big data is likely to yield higher returns when combined with complementary investments in managerial and workforce skills, reflecting the importance of capability constraints in the adoption of complex technologies. Second, Chapters 3 and 4 highlight that labour market policies affect firms through shocks of different scope, which shape the margins of adjustment. The 2009 compulsory schooling reform operated as a narrow labour-supply shock, affecting a relatively small segment of the workforce (underage workers in a subset of firms), and induced adjustment primarily through internal workforce reallocation rather than technological upgrading. In contrast, statutory minimum wage increases constituted a broad-based labour cost shock affecting a large share of workers, which created stronger incentives for firms to adopt labour-saving technologies. Third, these technological adjustments have important implications for workers, as the incidence of adjustment is uneven within firms. Chapter 4 shows that robot adoption is associated with a wage penalty among more exposed workers in adopting firms, with displacement among workers on fixed-term contracts, and with internal reallocation among permanent workers. This implies that labour market institutions shape the distribution of adjustment within firms across workers depending on their level of protection, highlighting the importance of accounting for these differences in policy design.

Four classes of limitations apply. First, measurement: robot adoption is identified from international trade records via CN code 847950, which captures direct imports but not domestically produced or indirectly acquired robots, so the adoption indicator is a lower bound; the occupation-based proxies in Chapters 3 and 4 capture workforce composition rather than equipment ownership. Second, scope: Chapters 3 and 4 are restricted to the Portuguese manufacturing sector, while Chapter 2 is multi-sectoral but constrained to firms in the IUTICE sample. Third, identification: the IUTICE waves for cloud computing and big data are intermittent, which limits panel-based inference in Chapter 2; in Chapter 3, the baseline ICT coefficient is not robust to age-based placebos and trend-augmented specifications; in Chapter 4, the firm-level effects of robot adoption are confounded by selection, which is why the worker-level analysis is

centred on within-firm variation. Fourth, external validity: the findings are specific to Portugal, a Southern European economy with a segmented labour market, a late digital-transition profile, and a binding statutory minimum wage; extending these findings to different institutional settings requires an appropriate mapping of the scale and nature of the shock.

Several extensions are natural. For Chapter 2, later IUTICE waves cover artificial intelligence and additional ICT tools, which would allow a study of multi-technology cohorts and of whether first-mover advantages in big data carry into AI adoption. For Chapters 3 and 4, a natural extension would be to refine the measurement of robotisation using more granular indicators that do not rely exclusively on international trade data. In Chapter 4, the mediated effect of minimum wage increases on worker outcomes via robots is not estimated in this thesis and would require an identification strategy that links the first-stage exposure regression and the within-firm second stage.

The thesis shows that firms do not respond to labour market pressures through a single uniform channel. The mediators – capability, technology complexity, institutional segmentation, and the scale of the shock – can be identified and are consistent across the three settings studied. Automation is therefore neither an inevitable consequence of institutional change nor an exogenous technological tide: it is a firm decision whose conditions, timing and distributional incidence can be disciplined empirically.

Bibliography

Aaronson, D., & Phelan, B. J. (2019). Wage shocks and the technological substitution of low-wage jobs. *The Economic Journal*, *129*(617), 1–34.

Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *The quarterly journal of economics*, *113*(4), 1055–1089.

Acemoglu, D. (2002). Directed technical change. *The review of economic studies*, *69*(4), 781–809.

Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy*, *118*(6), 1037–1078.

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics* (pp. 1043–1171, Vol. 4). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H., & Price, B. (2014). Return of the solow paradox? IT, productivity, and employment in us manufacturing. *American Economic Review*, *104*(5), 394–399.

Acemoglu, D., Koster, H. R. A., & Ozgen, C. (2023). *Robots and workers: Evidence from the Netherlands* (Working Paper No. 31009). National Bureau of Economic Research. https://doi.org/10.3386/w31009

Acemoglu, D., Lelarge, C., & Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. *AEA Papers and Proceedings*, *110*, 383–388.

Acemoglu, D., & Restrepo, P. (2018a). Artificial intelligence, automation, and work. In *The economics of artificial intelligence: An agenda* (pp. 197–236). University of Chicago Press.

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

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

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, *128*(6), 2188–2244.

Acemoglu, D., & Restrepo, P. (2022a). Demographics and automation. *The Review of Economic Studies*, *89*(1), 1–44.

Acemoglu, D., & Restrepo, P. (2022b). Tasks, automation, and the rise in US wage inequality. *Econometrica*, *90*(5), 1973–2016.

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, *83*(6), 2411–2451. https://doi.org/https://doi.org/10.3982/ECTA13408

Adachi, D., Kawaguchi, D., & Saito, Y. U. (2024). Robots and employment: Evidence from Japan, 1978–2017. *Journal of Labor Economics*, *42*(2), 591–634.

Aghion, P., Antonin, C., Bunel, S., & Jaravel, X. (2020). *What are the labor and product market effects of automation?: New evidence from France* (J. Miron, Ed.) (Cato Institute Research Briefs in Economic Policy. http://www.jstor.org/stable/resrep51923). Cato Institute. http://www.jstor.org/stable/resrep51923

Alexandre, F., Bação, P., Cerejeira, J., Costa, H., & Portela, M. (2022). Minimum wage and financially distressed firms: Another one bites the dust. *Labour Economics*, *74*, 102088.

Alexandre, F., Cerqueira, C., & Portela, M. (2026). Timing of cloud computing and big data adoption: Cohort characteristics and firm performance. *Economics of Innovation and New Technology*, 1–38.

Alexandre, F., Cruz, S., & Portela, M. (2021). Financial distress and the role of management in micro and small-sized firms [OECD Productivity Working Papers No. 2, OECD Publishing, Paris]. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3704119

Amador, J. (2024). Robots in Portuguese firms. *Banco de Portugal Economic Studies*, *X*(2), 3–24.

Amador, J., & Silva, C. (2023). The impact of ICT adoption on productivity: Evidence from Portuguese firm-level data [Working Papers No. 07/2023, Banco de Portugal].

Andrews, D., Criscuolo, C., & Gal, P. N. (2015). Frontier firms, technology diffusion and public policy: Micro evidence from OECD countries [OECD Productivity Working Papers No. 2, OECD Publishing, Paris]. https://doi.org/https://doi.org/10.1787/5jrql2q2jj7b-en

Andrews, D., Criscuolo, C., & Gal, P. N. (2016). The best versus the rest: The global productivity slowdown, divergence across firms and the role of public policy [OECD Productivity Working Papers No. 5, OECD Publishing, Paris]. https://doi.org/https://doi.org/10.1787/63629cc9-en

Angrist, J. D., & Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *Quarterly Journal of Economics*, *106*(4), 979–1014.

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. https://doi.org/10.1257/jep.29.3.3

Autor, D. H. (2019). Work of the past, work of the future. *AEA Papers and Proceedings*, *109*, 1–32. https://doi.org/10.1257/pandp.20191110

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, *103*(5), 1553–1597.

Autor, D. H., Dorn, D., & Hanson, G. H. (2015). Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, *125*(584), 621–646.

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

Barbosa, N. (2024). Artificial intelligence and exporting performance: Firm-level evidence from Portugal [Gabinete de Estratégia e Estudos, Ministério da Economia, Portugal]. *GEE Papers*, (183). https://www.gee.gov.pt/RePEc/WorkingPapers/GEE_PAPERS_183.pdf

Barbosa, N., & Faria, A. P. (2022). Digital adoption and productivity: Understanding micro drivers of the aggregate effect [Gabinete de Estratégia e Estudos, Ministério da Economia, Portugal]. *GEE Papers*, (162). https://www.gee.gov.pt/RePEc/WorkingPapers/GEE_PAPERS_162.pdf

Bartel, A., Ichniowski, C., & 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.

Bartelsman, E. J. (2013). ICT, reallocation and productivity [European Economy Economic Papers No. 486, European Commission].

Bastos, P., Flach, L., & Keller, K. (2023). *Robotizing to compete? Firm-level evidence* [Discussion Paper No. 467, LMU Munich and Humboldt University Berlin, Germany]. https://hdl.handle.net/10419/282158

Benmelech, E., & Zator, M. (2025). Robots and firm investment. *Journal of Financial Economics*. https://doi.org/10.1016/j.jfineco.2025.104054

Berlingieri, G., Calligaris, S., Criscuolo, C., & Verlhac, R. (2020). Laggard firms, technology diffusion and its structural and policy determinants [OECD Science, Technology and Industry Policy Papers No. 86, OECD Publishing, Paris]. https://doi.org/https://doi.org/10.1787/281bd7a9-en

Bessen, J., Goos, M., Salomons, A., & Van den Berge, W. (2025). What happens to workers at firms that automate? *Review of Economics and Statistics*, *107*(1), 125–141.

Bessen, J., Goos, M., Salomons, A., & van den Berge, W. (2020). Firm-level automation: Evidence from the Netherlands. *AEA Papers and Proceedings*, *110*, 389–393.

Blanas, S., Gancia, G., & Lee, S. Y. (2019). Who is afraid of machines? *Economic Policy*, *34*(100), 627–690.

Bloom, N., Brynjolfsson, E., Foster, L., Jarmin, R., Patnaik, M., Saporta-Eksten, I., & Van Reenen, J. (2019). What drives differences in management practices? *American Economic Review*, *109*(5), 1648–83. https://www.aeaweb.org/articles?id=10.1257/aer.20170491

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., & Roberts, J. (2012). Does management matter? Evidence from India. *The Quarterly Journal of Economics*, *128*(1), 1–51. https://doi.org/10.1093/qje/qjs044

Bloom, N., & Reenen, J. V. (2011). Human resource management and productivity. In D. Card & O. Ashenfelter (Eds.), *Handbook of labor economics* (pp. 1697–1767, Vol. 4, Part B). Elsevier. https://doi.org/10.1016/S0169-7218(11)02417-8

Bloom, N., Sadun, R., & Reenen, J. V. (2012). Americans do IT better: Us multinationals and the productivity miracle. *American Economic Review*, *102*(1), 167–201.

Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, *122*(4), 1351–1408. https://doi.org/10.1162/qjec.2007.122.4.1351

Boeri, T., & Garibaldi, P. (2007). Two tier reforms of employment protection: A honeymoon effect? *The Economic Journal*, *117*(521), 357–385.

Bonfiglioli, A., Crino, R., Fadinger, H., & Gancia, G. (2024). Robot imports and firm-level outcomes. *The Economic Journal*, *134*(664), 3428–3444.

Borowiecki, M., Pareliussen, J., Glocker, D., Kim, E. J., Polder, M., & Rud, I. (2021). The impact of digitalisation on productivity: Firm-level evidence from the netherlands. *OECD Economics Department Working Papers No. 1680, OECD Publishing, Paris*.

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, *117*(1), 339–376.

Brynjolfsson, E., Jin, W., & McElheran, K. (2021). The power of prediction: Predictive analytics, workplace complements, and business performance. *Business Economics*, *56*, 217–239.

Brynjolfsson, E., Li, J. F., Miranda, J., Seamans, R., & Wang, A. J. (2026, February). *Minimum wages and rise of the robots* (Working Paper No. 34895). National Bureau of Economic Research. http://www.nber.org/papers/w34895

Brynjolfsson, E., & McAfee, A. (2012). Winning the race with ever-smarter machines. *MIT Sloan Management Review*, *53*(2), 53–60. https://www.scopus.com/inward/record.uri?eid=2-s2.0-84856850008&partnerID=40&md5=57a29db6332672416d47f6efde861ad9

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 23–57). University of Chicago Press.

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, *13*(1), 333–72.

Busso, M., DiNardo, J., & McCrary, J. (2014). New evidence on the finite sample properties of propensity score reweighting and matching estimators. *Review of Economics and Statistics*, *96*(5), 885–897.

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.

Calvino, F., DeSantis, S., Desnoyers-James, I., Formai, S., Goretti, I., Lombardi, S., Manaresi, F., & Perani, G. (2022). Closing the italian digital gap: The role of skills, intangibles and policies [OECD Science, Technology and Industry Policy Papers No. 2022/126, OECD Publishing, Paris]. https://doi.org/https://doi.org/10.1787/e33c281e-en

Calvino, F., & Fontanelli, L. (2023). A portrait of AI adopters across countries [OECD Science, Technology and Industry Working Papers No. 2023/02, OECD Publishing, Paris]. https://doi.org/https://doi.org/https://doi.org/10.1787/0fb79bb9-en

Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: Methods and applications*. Cambridge University Press, Cambridge: New York.

Campos, M. M., & Reis, H. (2017). *Revisiting the returns to schooling in the Portuguese economy* (tech. rep.). Banco de Portugal.

Carneiro, P., Liu, K., & Salvanes, K. G. (2023). The supply of skill and endogenous technical change: Evidence from a college expansion reform. *Journal of the European Economic Association*, *21*(1), 48–92. https://doi.org/10.1093/jeea/jvac032

Carpenter, M. A. (2002). The implications of strategy and social context for the relationship between top management team heterogeneity and firm performance. *Strategic Management Journal*, *23*(3), 275–284. https://doi.org/https://doi.org/10.1002/smj.226

Casas, P., Castro Silva, H., Ribeiro, A. S., & Baptista, R. (2025, March). *Automação e inteligência artificial no mercado de trabalho português*. Fundação Francisco Manuel dos Santos, Lisboa, Portugal. https://ffms.pt/pt-pt/estudos/policy-papers

Cengiz, D., Dubé, A., Lindner, A., & Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, *134*(3), 1405–1454.

Centeno, M., & Novo, Á. A. (2012). Excess worker turnover and fixed-term contracts: Causal evidence in a two-tier system. *Labour Economics*, *19*(3), 320–328. https://doi.org/10.1016/j.labeco.2012.02.006

Cette, G., Nevoux, S., & Py, L. (2022). The impact of ICTs and digitalization on productivity and labor share: Evidence from French firms. *Economics of Innovation and New Technology*, *31*(8), 669–692.

Cette, G., & Spiezia, V. (2022). *Ict contribution to growth measurement: Do we miss a lot?* Edward Elgar Publishing.

Chen, J., & Roth, J. (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, *139*(2), 891–936. https://doi.org/10.1093/qje/qjad054

Cheng, H., Jia, R., Li, D., & Li, H. (2019). The rise of robots in China. *Journal of Economic Perspectives*, *33*(2), 71–88.

Cho, J., DeStefano, T., Kim, H., Kim, I., & Paik, J. H. (2023). What's driving the diffusion of next-generation digital technologies? *Technovation*, *119*, 102477.

Clemens, J., Kahn, L. B., & Meer, J. (2021). Dropouts need not apply? the minimum wage and skill upgrading. *Journal of Labor Economics*, *39*(S1), S107–S149.

Clemens, M. A., Lewis, E. G., & Postel, H. M. (2018). Immigration restrictions as active labor market policy: Evidence from the Mexican Bracero exclusion. *American Economic Review*, *108*(6), 1468–1487. https://doi.org/10.1257/aer.20170765

Conti, R., de Matos, M. G., & Valentini, G. (2023). Big data analytics, firm size, and performance. *Strategy Science*, *9*(2). https://doi.org/https://doi.org/10.1287/stsc.2022.0007

Cooper, R. W., & Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, *73*(3), 611–633. https://doi.org/10.1111/j.1467-937X.2006.00389.x

Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, *211*, 188–205.

Dai, X., & Qiu, Y. (2024). Minimum wage hikes and technology adoption: Evidence from u.s. establishments. *Journal of Financial and Quantitative Analysis*, *59*(4), 1620–1658.

Danzer, A. M., Feuerbaum, C., & Gaessler, F. (2024). Labor supply and automation innovation: Evidence from an allocation policy. *Journal of Public Economics*, *235*, 105136. https://doi.org/10.1016/j.jpubeco.2024.105136

Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, *19*(6), 3104–3153.

Daveri, F., & Parisi, M. L. (2015). Experience, innovation, and productivity: Empirical evidence from Italy's slowdown. *ILR Review*, *68*(4), 889–915. https://doi.org/10.1177/0019793915586971

Dayıoğlu, M., & Kırdar, M. G. (2022). Keeping kids in school and out of work: Compulsory schooling and child labor in turkey. *Journal of Human Capital*, *16*(4), 526–555.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, *110*(9), 2964–2996.

DeStefano, T., Kneller, R., & Timmis, J. (2023). Cloud computing and firm growth [*Forthcoming*]. *Review of Economics and Statistics*.

Devereux, P. J., & Hart, R. A. (2010). Forced to be rich? returns to compulsory schooling in britain. *The Economic Journal*, *120*(549), 1345–1364.

DGEEC. (2023). 75 anos de estatísticas da educação em Portugal, vol. III: Indicator I.1.0 – taxa de escolarização por nível de ensino e ano letivo [Direct data pages: https://info.dgeec.medu.pt/75anos-estatisticas-educacao-portugal/862/ and https://info.dgeec.medu.pt/75anos-estatisticas-educacao-portugal/863/. Accessed March 2026.].

DGEEC. (2024). Educação em números – Portugal 2024, table 1.1.5: Taxa real de escolarização (%) [Direct PDF link: https://www.dgeec.medu.pt/api/ficheiros/66a3a4676d6fdaac73b574dd. Accessed March 2026.].

Dinlersoz, E., & Wolf, Z. (2024). Automation, labor share, and productivity: Plant-level evidence from US manufacturing. *Economics of Innovation and New Technology*, *33*(4), 604–626.

Direção-Geral de Estatísticas da Educação e Ciência (DGEEC). (2026). Alunos matriculados no ensino secundário por nuts ii, programas orientados para jovens [Direct legacy table used for 2000/01–2002/03: https://estatisticas-educacao.dgeec.medu.pt/DireitoaEducacao/i1/I10_3.htm. Direct archive listing endpoint used to retrieve the annual workbook download URLs for 2003/04–2023/24: https://www.dgeec.medu.pt/api/postsano/652d2759121f641a986cc636. Example direct workbook link used in the rebuilt series (2022/23): https://www.dgeec.medu.pt/api/ficheiros/6690f4861638429280fe59fe. Accessed March 2026.].

Dixon, J., Hong, B., & Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, *67*(9), 5586–5605.

Draca, M., Machin, S., & Van Reenen, J. (2011). Minimum wages and firm profitability. *American Economic Journal: Applied Economics*, *3*(1), 129–151.

Duso, T., & Schiersch, A. (2022). Let's switch to the cloud: Cloud adaption and its effect on it investment and productivity. *DIW Berlin Discussion Paper No. 2017 - 33 S.*

Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M., & vom Berge, P. (2022). Reallocation effects of the minimum wage. *The Quarterly Journal of Economics*, *137*(1), 267–328. https://doi.org/10.1093/qje/qjab028

Engelstätter, B. (2012). It is not all about performance gains–enterprise software and innovations. *Economics of Innovation and New Technology*, *21*(3), 223–245.

Eurostat. (2024). Early leavers from education and training [Accessed March 2026].

Fan, H., Hu, Y., & Tang, L. (2021). Labor costs and the adoption of robots in China. *Journal of Economic Behavior and Organization*, *186*, 608–631.

Fernald, J., Inklaar, R., & Ruzic, D. (2024). The productivity slowdown in advanced economies: Common shocks or common trends? [*Forthcoming*]. *The Review of Income and Wealth*.

Freeman, R. B., Liu, X., Liu, Z., Song, R., & Xiong, R. (2025). The cause and consequence of robot adoption in China: Minimum wages and firms' responses. *Fundamental Research*, *5*(4), 1759–1770.

Gal, P. G., Nicoletti, G., Renault, T., Sorbe, S., & Timiliotis, C. (2019). Digitalisation and productivity: In search of the holy grail – firm-level empirical evidence from eu countries [OECD Economics Department Working Papers No. 1533, OECD Publishing, Paris]. https://doi.org/10.1787/5080f4b6-en

Giorcelli, M. (2019). The long-term effects of management and technology transfers. *American Economic Review*, *109*(1), 121–52. https://www.aeaweb.org/articles?id=10.1257/aer.20170619

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

Goulart, P., & Bedi, A. S. (2017). The evolution of child labor in portugal, 1850–2001. *Social Science History*, *41*(2), 227–254.

Graetz, G., & Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, *100*(5), 753–768.

Guadalupe, M., Kuzmina, O., & Thomas, C. (2012). Innovation and foreign ownership. *American Economic Review*, *102*(7), 3594–3627. https://doi.org/10.1257/aer.102.7.3594

Hall, B. H., Lotti, F., & Mairesse, J. (2013). Evidence on the impact of R&D and ICT investments on innovation and productivity in Italian firms. *Economics of Innovation and New Technology*, *22*(3), 300–328.

Haller, S. A., & Siedschlag, I. (2011). Determinants of ict adoption: Evidence from firm-level data. *Applied Economics*, *43*(26), 3775–3788.

Harasztosi, P., & Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, *109*(8), 2693–2727. https://doi.org/10.1257/aer.20171445

Hau, H., Huang, Y., & Wang, G. (2020). Firm response to competitive shocks: Evidence from China's minimum wage policy. *The Review of Economic Studies*, *87*(6), 2639–2671.

Heeringa, S. G., West, B. T., & Berglund, P. A. (2017). *Applied survey data analysis (2nd edn.)* Taylor; Francis, London. https://doi.org/https://doi.org/10.1201/9781315153278

Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, *108*(7), 1737–1772. https://doi.org/10.1257/aer.20161570

Humlum, A. (2021). *Robot adoption and labor market dynamics* [Mimeo, University of Chicago. PDF: https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/61a4e7d3dfd0aa67849b9918/1638197204858/humlumJMP.pdf]. https://static1.squarespace.com/static/5d35e72fcff15f0001b48fc2/t/61a4e7d3dfd0aa67849b9918/1638197204858/humlumJMP.pdf

Humlum, A. (2022). *Robot adoption and labor market dynamics* (tech. rep. No. 175). Rockwool Foundation Research Unit.

INE. (2021a). *Information and Communication Technologies Usage and e-commerce in Enterprises (Inquérito à Utilização das Tecnologias da Informação e Comunicação nas Empresas – IUTICE)* (Data) (Access to the original data through DGEEC (https://acreditacao.dgeec.medu.pt/)). Instituto Nacional de Estatística. Lisbon.

INE. (2021b). *Integrated Business Accounts System (Sistema de Contas Integradas das Empresas – SCIE)* (Data) (Access to the original data through DGEEC (https://acreditacao.dgeec.medu.pt/)). Instituto Nacional de Estatística. Lisbon.

INE. (2021c). *Personnel Records (Quadros de Pessoal – QP)* (Data) (Access to the original data through DGEEC (https://acreditacao.dgeec.medu.pt/)). Instituto Nacional de Estatística. Lisbon.

INE, Statistics Portugal. (2025). *Comércio internacional*. Instituto Nacional de Estatística, Lisboa.

Instituto Nacional de Estatística (INE). (2026). População residente por nuts ii, sexo e idade (ano a ano), indicador 0000256 [Metadata link: https://www.ine.pt/ine/json_indicador/pindicaMeta.jsp?varcd=0000256&lang=PT. Direct JSON endpoint for the full indicator payload used in the rebuild: https://www.ine.pt/ine/json_indicador/pindica.jsp?op=2&varcd=0000256&lang=PT. The figure denominator sums ages 15–18 for mainland NUTS II regions from that download. Accessed March 2026.].

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182.

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

Leon-Bravo, E. (2022). Does compulsory schooling impact labour market outcomes? evidence from the 1993 educational reform in mexico. (wp n.º 2).

Leone, F. (2024). *Global robots* (Working Paper). Bank of Italy and CESifo.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, *126*(2), 1029–1069.

Lordan, G., & Neumark, D. (2018). People versus machines: The impact of minimum wages on automatable jobs. *Labour Economics*, *52*, 40–53. https://doi.org/10.1016/j.labeco.2018.03.006

Lourenço, J. M. (2021). *The NOVAthesis LaTeX Template User's Manual*. NOVA University Lisbon. https://github.com/joaomlourenco/novathesis/raw/main/template.pdf

Mann, K., & Pozzoli, D. (2024). Robots and immigration. *Journal of Economic Behavior & Organization*, *227*. https://doi.org/10.1016/j.jebo.2024.106662

Montobbio, F., Staccioli, J., Virgillito, M. E., & Vivarelli, M. (2024). Labour-saving automation: A direct measure of occupational exposure. *The World Economy*, *47*(1), 332–361.

Müller, O., Fay, M., & Vom Brocke, J. (2018). The effect of big data and analytics on firm performance: An econometric analysis considering industry characteristics. *Journal of Management Information Systems*, *35*(2), 488–509.

Niebel, T., Rasel, F., & Viete, S. (2019). BIG data–BIG gains? Understanding the link between big data analytics and innovation. *Economics of Innovation and New Technology*, *28*(3), 296–316.

OECD. (2010). *Education at a glance 2010: OECD indicators* (tech. rep.). OECD Publishing. Paris. https://doi.org/10.1787/eag-2010-en

Oreopoulos, P. (2006). Estimating average and local average treatment effects of education when compulsory schooling laws really matter. *American Economic Review*, *96*(1), 152–175.

Pischke, J.-S., & von Wachter, T. (2008). Zero returns to compulsory schooling in Germany: Evidence and interpretation. *The Review of Economics and Statistics*, *90*(3), 592–598. https://doi.org/10.1162/rest.90.3.592

Reis, R. (2013). *The Portuguese slump and crash and the euro crisis* (tech. rep. No. 1). Brookings Institution.

San, S. (2023). Labor supply and directed technical change: Evidence from the termination of the bracero program in 1964. *American Economic Journal: Applied Economics*, *15*(1), 136–163.

Schivardi, F., & Schmitz, T. (2019). The IT revolution and Southern Europe's two lost decades. *Journal of the European Economic Association*, *18*(5), 2441–2486. https://doi.org/10.1093/jeea/jvz048

Segarra-Blasco, A., Teruel, M., & Cattaruzzo, S. (2022). Innovation, productivity and learning induced by export across European manufacturing firms. *Economics of Innovation and New Technology*, *31*(5), 387–415.

Stiroh, K. J. (2002). Information technology and the us productivity revival: What do the industry data say? *American Economic Review*, *92*(5), 1559–1576.

Stoneman, P., & Kwon, M. J. (1996). Technology adoption and firm profitability. *The Economic Journal*, *106*(437), 952–962. https://doi.org/10.2307/2235366

Strawinski, P., & Broniatowska, P. (2024). The impact of prolonging compulsory general education on the labour market. *Educational Studies*, *50*(4), 413–433.

Tiedtke, J. (2025, April). *Automation, firm size and skill groups* [LEM Working Paper 2025/09, Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy]. https://doi.org/10.57838/sssa/bz07-q629

Torres, S., Portugal, P., Addison, J. T., & Guimaraes, P. (2018). The sources of wage variation and the direction of assortative matching: Evidence from a three-way high-dimensional fixed effects regression model. *Labour Economics*, *54*, 47–60.

Van Ark, B. (2016). The productivity paradox of the new digital economy. *International Productivity Monitor*, *31*, 3–18.

Vieira, J. A. (1999). Returns to education in portugal. *Labour Economics*, *6*(4), 535–541.

Webb, M. (2020). *The impact of artificial intelligence on the labor market* [Unpublished manuscript. Available at https://www.michaelwebb.co/webb_ai.pdf]. https://www.michaelwebb.co/webb_ai.pdf

Wu, L., Hitt, L., & Lou, B. (2020). Data analytics, innovation, and firm productivity. *Management Science*, *66*(5), 2017–2039.

Wu, T.-H., Ting, P. J. L., Lin, M.-C., & Chang, C.-C. (2022). Corporate ownership and firm performance: A mediating role of innovation efficiency. *Economics of Innovation and New Technology*, *31*(4), 292–319.

A

Timing of cloud computing and big data adoption: Cohort characteristics and firm performance

A.1 Data and Variable Description

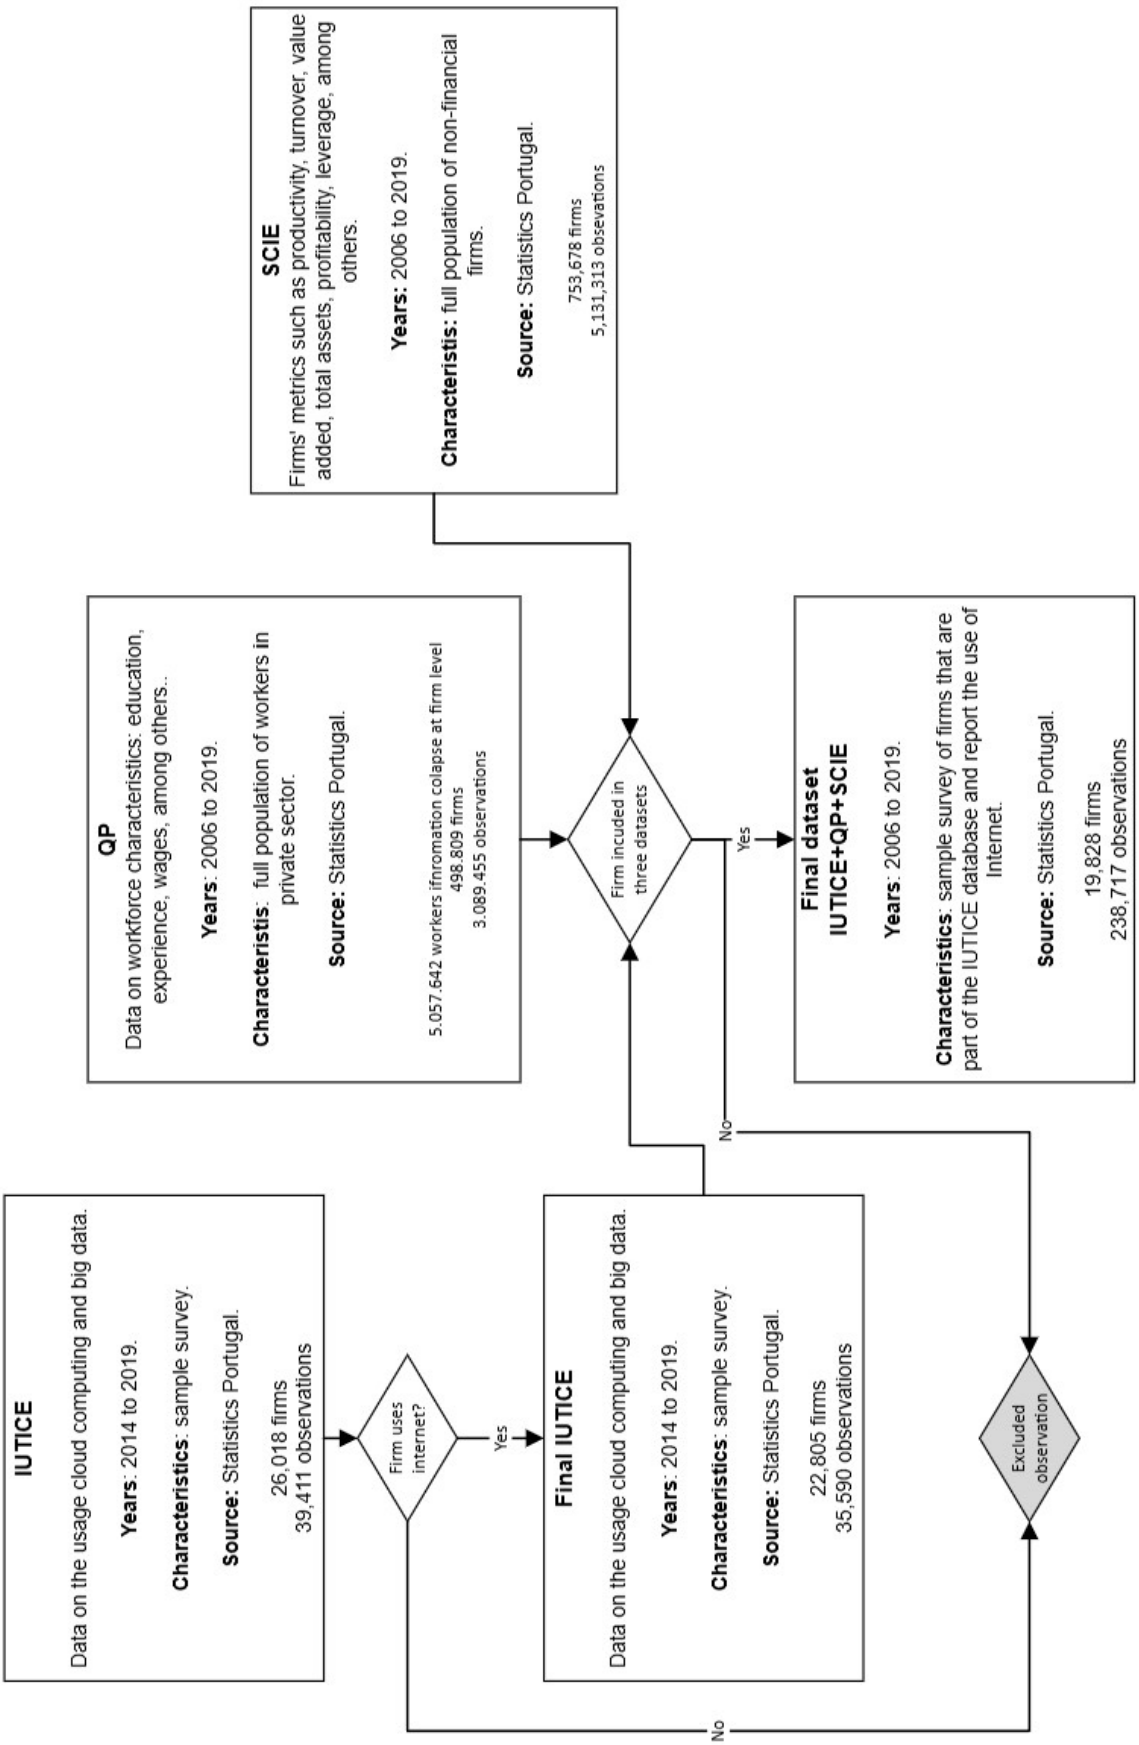


Figure 12: Flowchart illustrating the database construction process

Table 31: Variable description

Variable	Description	Source	Type
Average education of manager	Managers' team average years of education.	QP	EV
Previous management experience	Managers' team average years of experience as a man. in former firms.	QP	EV
Share High Educ Managers	The proportion of managers with a college diploma (%).	QP	EV
Average manager tenure	Managers' average tenure at the company.	QP	EV
Firm size	Micro, Small, Medium, and Large based on employees, total sales, and total assets.	SCIE	CV
Sector	Based on the 2-digit Classification of Portuguese Economic Activities.	SCIE	CV
Intermediate inputs	Accounting measure of intermediate inputs.	SCIE	CV
Capital stock	Accounting measure of capital stock.	SCIE	CV
Number of workers / Employment	Number of employees reported by the company each year.	SCIE	CV/DV
Share of higher-education workers	Employees (non-managers) with a college diploma (%).	QP	EV
Education of workforce	Workers' average years of education (excluding managers).	QP	EV
Labor productivity (LP)	$\frac{\text{Real value added}}{\text{Number of workers}}$	SCIE	DV
Gross value added (GVA)	Real gross value added.	SCIE	DV
Total factor productivity (TFP)	TFP was calculated as proposed by Ackerberg et al. (2015).	SCIE	DV
Profitability	$\frac{\text{EBIT}}{\text{Total Assets}}$	SCIE	CV
Capital intensity	$\frac{\text{Real Fixed Assets}}{\text{Number of workers}}$	SCIE	CV
Average wages	$\frac{\text{Real Labor Costs}}{\text{Number of Workers}}$	SCIE	CV
Leverage	$\frac{\text{Total Liabilities}}{\text{Total Assets}}$	SCIE	CV
Turnover	Real volume of sales and services provided.	SCIE	EV
Exporter	Dummy = 1 if firm exports $\geq 10\%$ of total sales for 3 consecutive years.	SCIE	CV
Importer	Dummy = 1 if firm imports $\geq 10\%$ of total purchases for 3 consecutive years.	SCIE	CV
Foreign owned	Dummy = 1 if firm is >50% foreign-owned, 0 otherwise.	QP	CV
Firm age	The number of years since the year of the constitution.	QP	CV
Investment in fixed assets	$\frac{\text{Investment in Fixed Assets}}{\text{Total Assets}}$	SCIE	CV
Share of computer users	$\frac{\text{Number of Computer Users}}{\text{Number of Workers}}$	SCIE and IUTICE	CV
Young firm	Dummy = 1 if firm is ≤10 years old 0 otherwise.	QP	CV
Cloud computing	Dummy = 1 if firm reports using cloud computing in the year analyzed.	IUTICE	DV
Big data	Dummy = 1 if firm reports using big data in the year analyzed.	IUTICE	DV

Notes: SCIE stands for Integrated Business Accounts System, QP stands for Personnel Records, and IUTICE stands for Information and Communication Technologies usage and e-commerce in enterprises. IV stands for Independent Variable, DV stands for Dependent Variable, and CV stands for Control Variable.

Table 32: Share of technological adopters by economic sector (%)

	2014	2016	2017	2018
	Panel A: Cloud Computing			
Manufacturing industries	4.8	6.7	10.3	11.2
Electricity, gas, steam and cold air	15.4	27.2	25.2	12.2
Water; sanitation, waste management and depollution	26.8	19.0	24.0	30.5
Construction	4.7	5.6	9.8	10.7
Wholesale and retail; car and motorcycle repair	6.0	9.7	9.3	11.7
Transport and storage	5.5	6.3	9.2	16.1
Accommodation and catering activities	6.5	8.7	8.7	9.7
Information and communication services	26.9	36.3	44.3	46.0
Real estate activities	6.5	16.1	15.5	15.1
Consulting, scientific, technical and similar services	9.1	18.0	20.3	25.7
Administrative and support services activities	10.2	17.2	20.5	24.2
Other services: Repair of comput. commun. equip.	14.1	11.2	32.1	34.2
	Panel B: Big Data			
Manufacturing industries		7.7		7.4
Electricity, gas, steam and cold air		22.0		20.8
Water; sanitation, waste management and depollution		24.0		20.9
Construction		8.6		5.7
Wholesale and retail; car and motorcycle repair		9.7		7.6
Transport and storage		8.6		18.5
Accommodation and catering activities		13.1		7.0
Information and communication services		21.2		20.6
Real estate activities		8.4		7.5
Consulting, scientific, technical and similar services		7.5		10.5
Administrative and support services activities		13.6		7.7
Other services: Repair of comput. commun. equip.		6.6		11.5

Notes: We report the percentage of adopters of each technology among the groups of firms divided by economic sectors in panels A to B. All years of the survey are included. Sampling weights provided by the survey are used in this report. Economic activity is classified according to the one-letter aggregation of the Portuguese Classification of Economic Activities (CAE), Revision 3, which corresponds to ISIC Revision 4. *Source:*Authors' calculations based on IUTICE and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

Table 33: Distribution of technological adopters by sector (%)

	2014	2016	2017	2018
	Panel A: Cloud Computing			
Manufacturing industries	13.4	12.4	16.4	15.3
Electricity, gas, steam and cold air	0.3	0.2	0.2	0.1
Water; sanitation, waste management and depollution	1.6	0.7	0.7	0.8
Construction	8.9	6.4	9.9	9.0
Wholesale and retail; car and motorcycle repair	31.8	32.6	26.1	26.6
Transport and storage	4.8	3.1	4.0	5.8
Accommodation and catering activities	8.2	7.5	6.6	6.4
Information and communication services	9.3	8.9	8.9	8.0
Real estate activities	1.7	3.5	2.9	2.9
Consulting, scientific, technical and similar services	14.4	18.4	18.0	18.6
Administrative and support services activities	5.6	6.1	6.0	6.3
Other services: Repair of comput. commun. equip.	0.1	0.2	0.2	0.2
	Panel B: Big Data			
Manufacturing industries		15.3		17.1
Electricity, gas, steam and cold air		0.2		0.2
Water; sanitation, waste management and depollution		1.0		0.9
Construction		10.5		8.2
Wholesale and retail; car and motorcycle repair		35.2		29.3
Transport and storage		4.6		11.3
Accommodation and catering activities		12.2		7.9
Information and communication services		5.6		6.1
Real estate activities		2.0		2.4
Consulting, scientific, technical and similar services		8.2		12.9
Administrative and support services activities		5.2		3.4
Other services: Repair of comput. commun. equip.		0.1		0.1

Notes: We report the distribution of adopters of each technology among economic sectors in panels A to B. All years of the survey are included. Sampling weights provided by the survey are used in this report. Economic activity is classified according to the one-letter aggregation of the Portuguese Classification of Economic Activities (CAE), Revision 3, which corresponds to ISIC Revision 4. *Source:* Authors' calculations based on IUTICE and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

Table 34: Distribution of technological cloud computing by main economic sector (%)

Subsector	Share
Wholesale and retail, car and motorcycle repair	
Sale, maintenance and repair of motor vehicles and motorcycles	17.28
Wholesale (including agents), except of motor vehicles and motorcycles	48.10
Retail trade, except of motor vehicles and motorcycles	34.62

Notes: Sampling weights provided by the survey are used in this report. *Source:* Authors' calculations based on IUTICE and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

Table 35: Distribution of technological big data by main sector (%)

Subsector	Share
Wholesale and retail, car and motorcycle repair	
Sale, maintenance and repair of motor vehicles and motorcycles	15.47
Wholesale (including agents), except of motor vehicles and motorcycles	28.75
Retail trade, except of motor vehicles and motorcycles	55.78

Notes: Sampling weights provided by the survey are used in this report. *Source:* Authors' calculations based on IUTICE and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

A.2 Results Without Propensity Score Reweighting

Table 36: First movers versus non-adopters – Impact of Cloud Computing (without propensity score reweighting)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	GVA		Turnover		Employment		LP		TFP	
Panel A										
Education Manager$_{t-1}$	-0.0003	-0.0002	-0.0010***	-0.0009***	-0.0005*	-0.0005	0.0001	0.0003	0.0002	0.0003*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Cloud Computing FM$_{t-1}$	0.1514***	0.1780***	0.1008***	0.1222***	0.0851***	0.0938***	0.0830***	0.1106***	0.0446***	0.0789**
	(0.007)	(0.037)	(0.005)	(0.033)	(0.004)	(0.016)	(0.006)	(0.033)	(0.006)	(0.035)
CC$_{t-1}$ × EM$_{t-1}$		-0.0022		-0.0017		-0.0007		-0.0022		-0.0028
		(0.003)		(0.002)		(0.001)		(0.002)		(0.002)
Panel B										
Tenure Manager$_{t-1}$	0.0003**	0.0004**	-0.0002	-0.0002	0.0003***	0.0003***	0.0003**	0.0003**	-0.0002	-0.0001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Cloud Computing FM$_{t-1}$	0.1514***	0.1569***	0.1000***	0.1056***	0.0849***	0.0808***	0.0833***	0.0874***	0.0446***	0.0494***
	(0.007)	(0.011)	(0.005)	(0.007)	(0.004)	(0.006)	(0.006)	(0.009)	(0.006)	(0.008)
CC$_{t-1}$ × TM$_{t-1}$		-0.0005		-0.0005		0.0004		-0.0004		-0.0004
		(0.001)		(0.001)		(0.000)		(0.000)		(0.000)
Panel C										
Prev. Manager Exper$_{t-1}$	-0.0013	-0.0006	0.0008	0.0017	-0.0008	-0.0010	-0.0020	-0.0018	-0.0000	0.0007
	(0.003)	(0.003)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.003)	(0.002)	(0.002)
Cloud Computing FM$_{t-1}$	0.1512***	0.1536***	0.1002***	0.1032***	0.0847***	0.0841***	0.0831***	0.0840***	0.0447***	0.0472***
	(0.007)	(0.007)	(0.005)	(0.007)	(0.004)	(0.004)	(0.006)	(0.006)	(0.006)	(0.006)
CC$_{t-1}$ × PME$_{t-1}$		-0.0068		-0.0085		0.0020		-0.0025		-0.0069
		(0.006)		(0.008)		(0.003)		(0.004)		(0.004)
Panel D										
Education workers$_{t-1}$	-0.0011	-0.0014*	-0.0013**	-0.0013**	-0.0010**	-0.0010*	-0.0007	-0.0009**	-0.0006**	-0.0009***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Cloud Computing FM$_{t-1}$	0.1517***	0.1145***	0.1009***	0.1009***	0.0854***	0.0831***	0.0834***	0.0562***	0.0451***	0.0150
	(0.007)	(0.015)	(0.005)	(0.005)	(0.004)	(0.009)	(0.006)	(0.011)	(0.006)	(0.010)
CC$_{t-1}$ × EW$_{t-1}$		0.0034**				0.0002		0.0025**		0.0028**
		(0.002)				(0.001)		(0.001)		(0.001)

Notes: All dependent variables are in logs. LP stands for labor productivity. TFP stands for total factor productivity. Coefficients were obtained using sampling weights and bootstrap standard errors. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector and computer user share. For the labor productivity, GVA, and turnover models, we additionally control for capital stock, intermediate inputs, and number of employees. For the TFP and employment models, firm size is included as a categorical control. The total number of observations is 3,366. *Significance levels:* *** 1%; ** 5%; * 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE, 2021a, INE, 2021c, INE, 2021b).

Table 37: First movers versus followers – Long-term impact of Cloud Computing (without propensity score reweighting)

	(1)	(2)	(3)	(4)	(5)
	ΔGVA	ΔTurnover	ΔEmployment	ΔLP	ΔTFP
			Panel A		
Education manager	0.0435***	0.0391**	0.0194	0.0164	0.0096
	(0.015)	(0.019)	(0.014)	(0.011)	(0.009)
CC Adopter	-0.2660***	-0.3961***	-0.1282	-0.0750	-0.0718
	(0.092)	(0.115)	(0.080)	(0.082)	(0.057)
			Panel B		
Tenure manager	-0.0200***	-0.0148**	-0.0155***	-0.0034	-0.0023
	(0.005)	(0.006)	(0.005)	(0.004)	(0.004)
CC Adopter	-0.2574***	-0.3891***	-0.1444*	-0.0728	-0.0734
	(0.092)	(0.115)	(0.078)	(0.082)	(0.056)
			Panel C		
Prev. manager exper	0.0455	0.0553	0.0386	-0.0049	0.0273
	(0.046)	(0.037)	(0.033)	(0.031)	(0.026)
CC Adopter	-0.2574***	-0.3869***	-0.1226	-0.0741	-0.0683
	(0.098)	(0.115)	(0.082)	(0.081)	(0.057)
			Panel D		
Education Workers	0.0164	0.0223	0.0211	-0.0076	0.0066
	(0.026)	(0.022)	(0.022)	(0.016)	(0.012)
CC Adopter	-0.2587***	-0.3880***	-0.1212	-0.0751	-0.0696
	(0.094)	(0.112)	(0.078)	(0.080)	(0.056)

Notes: All dependent variables are in logs. LP stands for labor productivity. TFP stands for total factor productivity. The coefficients were obtained using the Long Differences Model 2013-2019. Sampling weights were used. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, and computer user share. In the labor productivity and turnover models, we also account for capital stock, intermediate inputs, and the number of employees. In the GVA model, we control for capital stock and the number of employees. Finally, we control for firm size in the labor productivity and employment models. Robust standard errors are in parenthesis. The total number of observations is 909. *Significance levels:* ***, 1%; **, 5%; *, 10%. *Source:* Authors' calculations based on IUTICE, QP and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

Table 38: First movers versus non-adopters – Impact of Big Data (without propensity score reweighting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	GVA		Turnover		Employment		LP		TFP	
					Panel A					
Education Manager$_{t-1}$	0.0011***	0.0012***	-0.0001	0.0000	0.0002	0.0004	0.0009***	0.0009***	0.0010***	0.0008***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Big Data FM$_{t-1}$	0.1256***	0.1360***	0.0875***	0.0983***	0.0747***	0.0947***	0.0684***	0.0671***	0.0337***	0.0168**
	(0.003)	(0.011)	(0.003)	(0.011)	(0.003)	(0.006)	(0.002)	(0.006)	(0.003)	(0.007)
BD$_{t-1}$ × EM$_{t-1}$		-0.0009		-0.0009		-0.0017***		0.0001		0.0014***
		(0.001)		(0.001)		(0.000)		(0.000)		(0.001)
					Panel B					
Tenure Manager$_{t-1}$	-0.0002	-0.0002	-0.0002	-0.0002	-0.0000	-0.0001	-0.0002*	-0.0002	-0.0001	-0.0001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Big Data FM$_{t-1}$	0.1262***	0.1243***	0.0873***	0.0889***	0.0748***	0.0688***	0.0689***	0.0695***	0.0344***	0.0432***
	(0.003)	(0.005)	(0.003)	(0.005)	(0.003)	(0.004)	(0.002)	(0.004)	(0.003)	(0.006)
BD$_{t-1}$ × TM$_{t-1}$		0.0002		-0.0001		0.0005		-0.0001		-0.0008**
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
					Panel C					
Prev. Manager Exper$_{t-1}$	-0.0003	-0.0002	0.0006	0.0007	0.0006	0.0007	-0.0006	-0.0006	-0.0009*	-0.0012**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
Big Data FM$_{t-1}$	0.1264***	0.1268***	0.0874***	0.0878***	0.0747***	0.0751***	0.0691***	0.0691***	0.0346***	0.0331***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
BD$_{t-1}$ × PME$_{t-1}$		-0.0006		-0.0006		-0.0006		-0.0000		0.0024
		(0.001)		(0.001)		(0.001)		(0.001)		(0.002)
					Panel D					
Education workers$_{t-1}$	0.0013***	0.0015***	-0.0008***	-0.0006**	-0.0001	0.0000	0.0013***	0.0014***	0.0011**	0.0010**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
Big Data FM$_{t-1}$	0.1262***	0.1472***	0.0875***	0.1029***	0.0748***	0.0818***	0.0689***	0.0748***	0.0343***	0.0271***
	(0.003)	(0.010)	(0.003)	(0.007)	(0.003)	(0.007)	(0.002)	(0.009)	(0.003)	(0.010)
BD$_{t-1}$ × EW$_{t-1}$)		-0.0021**		-0.0015**		-0.0007		-0.0006		0.0007
		(0.001)		(0.001)		(0.001)		(0.001)		(0.001)

Notes: All dependent variables are in logs. LP stands for labor productivity. TFP stands for total factor productivity. Coefficients were obtained using sampling weights and bootstrap standard errors. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector and computer user share. For the labor productivity, GVA, and turnover models, we additionally control for capital stock, intermediate inputs, and number of employees. For the TFP and employment models, firm size is included as a categorical control. The total number of observations is 3,891. *Significance levels:* *** 1%; ** 5%; * 10%. *Source:* Authors' calculations based on IUTICE, QP, and SCIE (INE, 2021a, INE, 2021c, INE, 2021b).

Table 39: First movers versus followers – Long-term impact of Big Data (without propensity score reweighting)

	(1)	(2)	(3)	(4)	(5)
	ΔGVA	ΔTurnover	ΔEmployment	ΔLP	ΔTFP
			Panel A		
Education manager	-0.0371***	-0.0105	-0.0059	-0.0300***	-0.0175
	(0.014)	(0.011)	(0.007)	(0.011)	(0.014)
Big data FM adopters	0.1616**	0.1703**	0.0188	0.1925***	0.0227
	(0.066)	(0.073)	(0.047)	(0.054)	(0.042)
			Panel B		
Tenure manager	-0.0032	-0.0136**	-0.0094***	0.0073*	0.0061
	(0.005)	(0.005)	(0.003)	(0.004)	(0.005)
Big data FM adopters	0.1214*	0.1286**	-0.0248	0.1837***	0.0323
	(0.064)	(0.062)	(0.042)	(0.056)	(0.046)
			Panel C		
Prev. manager exper	-0.0122	0.0287	-0.0037	-0.0092	-0.0048
	(0.043)	(0.046)	(0.037)	(0.015)	(0.013)
Big data FM adopters	0.1286**	0.1623**	0.0149	0.1658***	0.0107
	(0.064)	(0.068)	(0.046)	(0.054)	(0.046)
			Panel D		
Education workers	0.0166	0.0106	0.0254**	-0.0064	-0.0182
	(0.021)	(0.021)	(0.012)	(0.018)	(0.020)
Big data FM adopters	0.1233**	0.1574**	0.0121	0.1684***	0.0114
	(0.062)	(0.071)	(0.043)	(0.054)	(0.045)

Notes: All dependent variables are in logs. LP stands for labor productivity. TFP stands for total factor productivity. The coefficients were obtained using the Long Differences Model 2015-2019. Sampling weights were used. All models include controls for exporter activity, leverage, foreign ownership status, importer activity, firm age, sector and computer user share. In the case of labor productivity, GVA and turnover models, we also account for capital stock, intermediate inputs, and the number of employees. In the case of TFP and employment models, we also control for firm size. Robust standard errors are in parenthesis. The total number of observations is 773. *Significance levels:* ***, 1%; **, 5%; *, 10%. *Source:* Authors' calculations based on IUTICE, QP and SCIE (INE, 2021a, INE, 2021c and INE, 2021b).

B

Keeping teenagers in school and firm adjustment: evidence from Portuguese manufacturing firms

B.1 Additional descriptive statistics

This appendix reports additional descriptive statistics. Unless stated otherwise, the robot descriptives in Figure A1 and Tables A1–A4 use the strict Tier 1 baseline measure based on CN6 847950. Figure A1 shows the cumulative value of strict robot imports and the annual number of strict robot-importing firms. Tables A1–A3 provide additional firm-level descriptive statistics, and Table A4 reports the distribution of strict adopters across manufacturing sectors.

B.2 Robustness checks

This appendix reports additional sensitivity checks. Table B1 reports the pooled firm-level DiD benchmark without firm fixed effects, moved out of the main text because it is descriptive only and systematically more expansionary than the baseline fixed-effects estimates. Appendix Table B3 does the same for the exit margin: it reports the pooled logit benchmark, moved out of the main text because it lacks firm fixed effects and is used only as a descriptive comparison to Table 19. Table B4 retains the broader five-outcome block with the cumulative import measures, which is now referenced from the main text but no longer occupies the central table position. Table B5 adds a short transformation robustness check that replaces $\log(1+Y)$ with $\operatorname{asinh}(Y)$ for the two cumulative import outcomes only.

Tables B6 and B7 report an alternative control specification that retains firm, sector × year, and region × year fixed effects and baseline education interacted with year dummies, but omits baseline size × year. Table B8 keeps the baseline binary treatment and the same fixed effects but adds, in sequence, baseline firm age × year and baseline workforce average age × year to the baseline control set. We also re-estimate the baseline specification without the baseline education or baseline size interacted controls (Table B9), using a non-linear probit specification for the binary outcomes (Table B10), and on a shorter panel window (2005–2015) to assess sensitivity to the sample period (Table B11).

Table A1: Descriptive statistics for firms established before and after the 2009 reform

Variable	Statistic	Pre-reform	After reform
Robot (0/1)	p1	0.0000	0.0000
	Mean	0.0006	0.0012
	p50	0.0000	0.0000
	p99	0.0000	0.0000
Import robots (ths, 2024 euro)	p1	0.00	0.00
	Mean	0.12	0.70
	p50	0.00	0.00
	p99	0.00	0.00
Share of minors	p1	0.0000	0.0000
	Mean	0.0045	0.0006
	p50	0.0000	0.0000
	p99	0.1250	0.0000
Years of education	p1	4.00	4.00
	Mean	7.66	8.97
	p50	7.33	9.00
	p99	15.00	15.00
Growth sales	p1	-1.13	-1.19
	Mean	0.05	0.18
	p50	0.01	0.08
	p99	1.77	2.79
ICT Worker	p1	0.00	0.00
	Mean	0.02	0.02
	p50	0.00	0.00
	p99	1.00	1.00
Hourly wage (2024 euro)	p1	2.98	3.39
	Mean	4.86	5.28
	p50	4.08	4.47
	p99	15.57	17.41
Number workers	p1	1.00	1.00
	Mean	9.87	8.78
	p50	4.00	3.00
	p99	99.00	88.00
Sales (ths, 2024 euro)	p1	0.00	0.00
	Mean	1225.72	1084.73
	p50	164.67	168.55
	p99	16051.33	14462.43
Observations		251,372	98,872
Number of firms		36,775	20,761

Notes: Manufacturing firms only. The robot rows use the strict Tier 1 measure. The table uses the full 1996–2021 descriptive panel and divides firms into those established before the increase in mandatory schooling (2002 to 2008) and those established after it (2010 to 2019). These are descriptive cohort moments only and combine reform-era changes with broader secular shifts in manufacturing.
Source: Authors' calculations using QP, CI, and SCIE data, 1996–2021.

Table A2: Strict Tier 1 robot adoption rates by firm size, 1996–2021

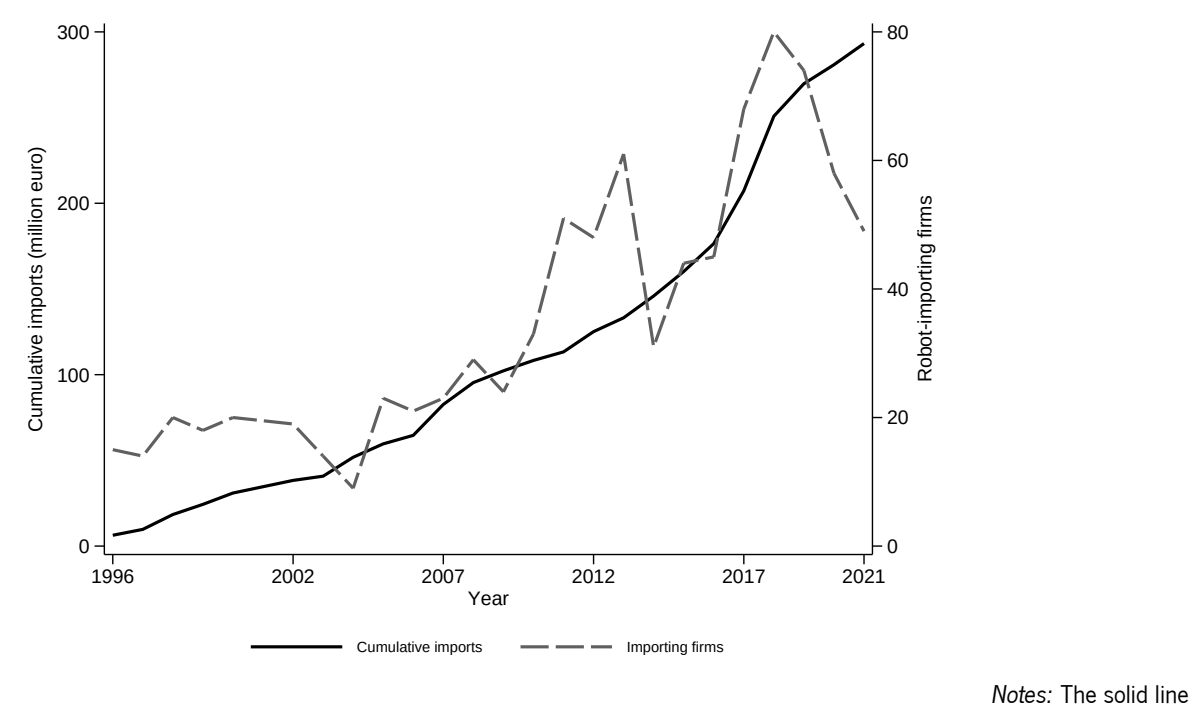
Year	Micro (%)	Small (%)	Medium (%)	Large (%)	Total (%)
1996	0.000	0.056	0.295	0.855	0.043
1997	0.004	0.052	0.162	1.187	0.037
1998	0.004	0.042	0.206	3.030	0.052
1999	0.004	0.030	0.364	1.497	0.044
2000	0.000	0.050	0.491	0.990	0.047
2002	0.000	0.040	0.404	2.069	0.045
2003	0.000	0.030	0.395	0.733	0.033
2004	0.003	0.019	0.176	0.690	0.021
2005	0.007	0.065	0.305	2.473	0.053
2006	0.000	0.067	0.398	1.859	0.049
2007	0.011	0.029	0.406	3.030	0.056
2008	0.011	0.040	0.590	3.462	0.072
2009	0.008	0.032	0.646	2.521	0.063
2010	0.004	0.127	0.682	3.376	0.097
2011	0.018	0.208	1.111	3.704	0.155
2012	0.014	0.157	0.947	6.780	0.154
2013	0.038	0.264	1.211	4.741	0.198
2014	0.010	0.077	0.639	4.803	0.100
2015	0.019	0.152	1.160	2.429	0.141
2016	0.024	0.113	0.970	4.781	0.143
2017	0.024	0.210	1.679	4.444	0.216
2018	0.038	0.172	1.724	7.971	0.255
2019	0.020	0.265	1.514	6.360	0.246
2020	0.025	0.258	1.189	3.226	0.195
2021	0.021	0.158	1.047	4.225	0.169

Notes: Rates are shown in percent for the strict Tier 1 importer measure; small positive values are retained to three decimals. *Source:* Authors' calculations using QP, CI, and SCIE data, 1996–2021.

Table A3: Strict Tier 1 robot adoption rates by manufacturing sector

Manufacturing activity group	1996	2002	2007	2012	2017	2021
Basic metals and metalworking	0.111	0.114	0.178	0.423	0.601	0.375
Chemicals, rubber and non-metallic	0.101	0.130	0.110	0.385	0.557	0.550
Electrical equipment and furniture	0.000	0.000	0.021	0.000	0.000	0.032
Food, beverages and tobacco	0.024	0.000	0.000	0.018	0.000	0.019
Textiles, apparel and leather	0.011	0.010	0.000	0.030	0.000	0.000
Wood, paper and printing	0.000	0.000	0.000	0.000	0.072	0.026

Notes: Entries are sector-specific adoption rates in percent for selected years using the strict Tier 1 importer measure. *Source:* Authors' calculations using QP, CI, and SCIE data, 1996–2021.



Notes: The solid line plots the cumulative value of strict Tier 1 robot imports (CN6 847950) across all manufacturing firms, in millions of 2024 euros (left axis). The dashed line plots the annual number of manufacturing firms recording at least one Tier 1 robot import (right axis). The year 2001 is absent because *Quadros de Pessoal* was not collected that year. *Source:* Authors' calculations using CI data for the manufacturing sample defined from QP, 1996–2021.

Figure A1: Strict robot import dynamics (Tier 1), 1996–2021

Table A4: Distribution of strict Tier 1 robot adopters across manufacturing sectors

Manufacturing activity group	1996	2002	2007	2012	2017	2021
Basic metals and metalworking	60.0	63.2	73.9	66.7	69.1	59.2
Chemicals, rubber and non-metallic	26.7	31.6	21.7	27.1	26.5	34.7
Electrical equipment and furniture	0.0	0.0	4.3	0.0	0.0	2.0
Food, beverages and tobacco	6.7	0.0	0.0	2.1	0.0	2.0
Textiles, apparel and leather	6.7	5.3	0.0	4.2	0.0	0.0
Wood, paper and printing	0.0	0.0	0.0	0.0	4.4	2.0

Notes: Entries report each sector's share of all strict Tier 1 robot-importing firms in the indicated year. *Source:* Authors' calculations using QP, CI, and SCIE data, 1996–2021.

Table B12 reports a narrow PPML sensitivity for annual robot importing and cumulative robot-import value in levels under the binary exposure only. It avoids the $\log(1 + Y)$ transformation for non-negative skewed outcomes and checks whether a level-based mean model changes the sign ranking of the results. It is not a like-for-like replacement for the baseline specification: the PPML implementation uses firm and common year effects rather than the full sector $\times$ year and region $\times$ year structure, and it drops firms with all-zero outcome paths under separation, leaving only a small retained sample. Because the treatment timing is common rather than staggered, Table B17 compares the baseline binary-exposure TWFE estimates with a single-cohort imputation ATT and with a trend-augmented TWFE specification that allows exposed firms to follow differential linear trends. Tables B18 and B19 then implement a design-based matching exercise that exact-matches on baseline sector, region, and size class, and uses three-nearest-neighbour matches on pre-2009 outcomes and firm characteristics to reweight the binary-exposure specification. Table B20 reports a comparison of firm-level versus sector–region cell-level clustered standard errors for the binary-exposure TWFE specification.

Supplementary Tables B13, B14, B15, and B16, together with the appendix figures, report the broader annual and cumulative outcome comparisons across the alternative CI-based, occupation-based, and threshold-based exposure definitions.

The unifying message is more conservative than the pooled and education-only benchmarks and broadly aligned with the baseline size-adjusted specification. Once firm fixed effects and baseline size-specific time paths are allowed for, the positive robot-import coefficients visible in the pooled benchmark and in the education-only control specification disappear or reverse sign under the binary, continuous-share, and percentile treatments, while ICT remains positive. In the main annual comparison, Tier 1 and Tier 2 imports remain near zero, the flexible occupation proxy is positive under all three main-text treatment codings, and ICT is the strongest annual margin. The threshold-based binary table makes the same point in a more demanding way: a 5% cutoff leaves a smaller positive ICT coefficient, but the 10% cutoff is largely null and does not revive the import margins. The education-only annual table shows the complementary contrast: once baseline size $\times$ year is omitted, binary robot coefficients turn positive again, which illustrates how sensitive the robot margin is to a control that is first-order given the exposure-size correlation in the data. Adding baseline firm age and workforce average age modestly attenuates ICT but leaves the annual robot margins null and makes the cumulative robot coefficients slightly more negative. Removing controls altogether leaves the qualitative ranking of outcomes broadly similar to the baseline table rather than to the education-only specification. The PPML estimates are positive but very imprecise and are obtained on a highly selected retained sample of roughly five thousand firm-year observations, so they do not strengthen the case for positive annual robot adoption or for positive cumulative robot-import value. Interpreted on the PPML scale, the binary coefficients correspond to multiplicative changes in the conditional mean: roughly 28–29% for the importer probability and about 7% for cumulative import value. But none of these coefficients is statistically distinguishable from zero, and none is estimated on a sample or fixed-effects structure comparable to the baseline TWFE specification. The broader automation tiers do not restore a positive cumulative-import pattern in the baseline binary TWFE specification, while the flexible

occupation proxy emerges as the most informative annual automation-related margin outside ICT. The single-cohort imputation ATT remains close to the baseline TWFE estimates, but trend-augmented TWFE and matched designs attenuate the coefficients sharply. Under sector–region clustering, ICT remains significant, whereas the import-based robot outcomes remain weak.

Table B1: Pooled firm-level DiD estimates without firm fixed effects

	Robot Dummy	Cumulative Robot imports (log)	TWFE Worker
Panel A: Post at 2009 and minors as treatment			
Minors	0.0005*	0.0526***	0.0308***
	(0.000)	(0.007)	(0.002)
[1em] Post 2009	0.0008***	0.0526***	0.0029***
	(0.000)	(0.003)	(0.001)
[1em] Minors x Post 2009	0.0013***	0.0734***	0.0123***
	(0.000)	(0.012)	(0.002)
Observations	428,090	428,090	428,090
Panel B: Post at 2009 and share of minors as treatment			
Share minors	-0.0048***	-0.2156***	-0.2052***
	(0.001)	(0.019)	(0.013)
[1em] Post 2009	0.0010***	0.0619***	0.0043***
	(0.000)	(0.003)	(0.001)
[1em] Share minors x Post 2009	-0.0003	-0.0660**	0.0664***
	(0.001)	(0.030)	(0.014)
Observations	428,090	428,090	428,090
Panel C: Post at 2009 and distribution share of minors as treatment			
Perc. share minors	-0.0000	0.0060***	0.0010**
	(0.000)	(0.002)	(0.000)
[1em] Post 2009	0.0009***	0.0575***	0.0034***
	(0.000)	(0.003)	(0.001)
[1em] Perc. share minors x Post 2009	0.0003**	0.0128***	0.0036***
	(0.000)	(0.003)	(0.001)
Observations	428,090	428,090	428,090

Table B2: Pooled firm-level DiD estimates without firm fixed effects (continued)

	Robot Dummy	Cumulative Robot imports (log)	TWFE Worker
Panel D: Post at 2012 and minors as treatment			
Minors	0.0006**	0.0587***	0.0315***
	(0.000)	(0.007)	(0.002)
[1em] Post 2012	0.0009***	0.0637***	0.0022***
	(0.000)	(0.004)	(0.001)
[1em] Minors x Post 2012	0.0015***	0.0837***	0.0178***
	(0.001)	(0.014)	(0.003)
Observations	353,217	353,217	353,217
Panel E: Post at 2012 and share of minors as treatment			
Share minors	-0.0048***	-0.2813***	-0.2177***
	(0.001)	(0.023)	(0.012)
[1em] Post 2012	0.0011***	0.0747***	0.0042***
	(0.000)	(0.004)	(0.001)
[1em] Share minors x Post 2012	-0.0024	-0.0928**	0.0710***
	(0.002)	(0.045)	(0.015)
Observations	353,217	353,217	353,217
Panel F: Post at 2012 and distribution share of minors as treatment			
Perc. share minors	0.0000	0.0055***	0.0008*
	(0.000)	(0.002)	(0.000)
[1em] Post 2012	0.0010***	0.0691***	0.0028***
	(0.000)	(0.004)	(0.001)
[1em] Perc. share minors x Post 2012	0.0003**	0.0157***	0.0053***
	(0.000)	(0.004)	(0.001)
Observations	353,217	353,217	353,217

Notes: Pooled difference-in-differences linear model without firm fixed effects, estimated with heteroskedasticity-robust standard errors. Columns report annual Tier 1 importer, cumulative Tier 1 imports $\log(1+\cdot)$, and TWFE worker. Panels A–C use the Post-2009 sample and report, respectively, the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment; Panels D–F repeat the same sequence on the Post-2012 sample. Within each panel, rows report the treatment main effect, the post indicator, and the treatment × post interaction. Controls include year dummies, sector indicators, region indicators, and baseline average years of education measured in 2008. Because firm fixed effects are omitted, this table is descriptive and not directly comparable to Table 16. *Source:* Authors' calculations using QP and CI data, 2002–2019.

Supplementary Tables B13, B14, and B15 report the broader annual and cumulative outcome comparisons directly. Under the baseline binary exposure, the annual import margins remain small and uneven across Tier 1, Tier 2, and Tier 3, while the cumulative log measures are negative or null rather than becoming more supportive as the import definition widens. The flexible occupation proxy is the only annual automation-related margin outside TWFE that turns clearly positive in the supplementary annual comparison, especially in Panel B. In the main annual comparison, annual Tier 1 and Tier 2 importing remain near zero, the flexible occupation proxy is positive, and TWFE is larger than the occupation margin. The corresponding event-study diagnostics are more useful for comparing design quality than for validating a stronger automation effect. Figure 5 reports the annual comparison across Tier 1, Tier 2,

Table B3: Pooled logit estimates for firm exit without firm fixed effects

	Market exit
Panel A: Post at 2009	
Minors x Post 2009	-0.0013
	(0.002)
Share minors x Post 2009	-0.0251
	(0.021)
Share minors (dist) x Post 2009	-0.0003
	(0.001)
Observations	279,185
Panel B: Post at 2012	
Minors x Post 2012	0.0027
	(0.002)
Share minors x Post 2012	-0.0026
	(0.024)
Share minors (dist) x Post 2012	0.0007
	(0.001)
Observations	243,383

Notes: Pooled logit model without firm fixed effects, estimated on the SCIE-linked exit sample with firm-clustered standard errors. Entries are average marginal effects. The outcome equals 1 in the firm's last observed year of SCIE activity, following the business-demography definition described in Section 4. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Rows report the same three treatment codings used in the main-text exit table: the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. Because firm fixed effects are omitted, this table is descriptive and not directly comparable in credibility to Table 19.
Source: Authors' calculations using SCIE and QP data, 2006–2019.

Tier 3, the two occupation proxies, and TWFE employment. On that margin, the strict Tier 1 importer remains the weakest timing case, while Tier 2 ($p = 0.289$), Tier 3 ($p = 0.603$), the occupation proxy (flexible) ($p = 0.560$), and TWFE ($p = 0.193$) do not reject the pre-trend null. The strict occupation proxy is also shown, but its current event-study specification does not support a comparable joint pre-trend test because the pre-reform interactions are absorbed or omitted. Appendix Figure B1 then reports the cumulative event studies for Tier 1, Tier 2, and Tier 3, none of which rejects the pre-trend null ($p = 0.493$, $p = 0.537$, and $p = 0.737$). Figure 4 and Appendix Figures B3 and B4 show that broadening the annual measure reduces some zero-driven jumps, but the relative ranking of more exposed and less exposed firms remains sensitive to the treatment cutoff. We therefore keep these broader measures as additional sensitivity checks rather than as replacements for the main annual comparison in Table 16.

Table B4: Five-outcome firm-level DiD estimates with annual and cumulative import outcomes

	Tier 1 importer	Tier 2 importer	ICT worker	Tier 1 cum. imports (log)	Tier 2 cum. imports (log)
Panel A: Post 2009					
Minors × Post 2009	-0.0002	-0.0000	0.0116***	-0.0400*	-0.0508**
	(0.001)	(0.001)	(0.004)	(0.020)	(0.025)
Share minors × Post 2009	-0.0019	-0.0025	0.0337**	-0.1615**	-0.3398***
	(0.002)	(0.003)	(0.017)	(0.065)	(0.088)
Percentile share minors × Post 2009	-0.0001	-0.0002	0.0037***	-0.0153***	-0.0254***
	(0.000)	(0.000)	(0.001)	(0.005)	(0.007)
Observations	428,076	428,076	428,076	428,076	428,076
Panel B: Post 2012					
Minors × Post 2012	-0.0003	0.0003	0.0163***	-0.0421*	-0.0466
	(0.001)	(0.001)	(0.004)	(0.024)	(0.028)
Share minors × Post 2012	-0.0035*	-0.0031	0.0435**	-0.1447*	-0.3098***
	(0.002)	(0.003)	(0.019)	(0.082)	(0.106)
Percentile share minors × Post 2012	-0.0002	-0.0001	0.0047***	-0.0154**	-0.0238***
	(0.000)	(0.000)	(0.001)	(0.006)	(0.008)
Observations	353,203	353,203	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes
Region × Year FE	Yes	Yes	Yes	Yes	Yes
Baseline educ./size × year	Yes	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual Tier 1 importer, annual Tier 2 importer, TWFE worker, cumulative Tier 1 imports $\log(1 + \cdot)$, and cumulative Tier 2 imports $\log(1 + \cdot)$. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. The cumulative columns are included because robot purchases are lumpy and may affect accumulated robot-related import stocks even when the annual importer margin is too thin to move cleanly; they should therefore be read as stock-like supporting outcomes rather than direct annual adoption effects. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B5: Cumulative import outcomes under the inverse hyperbolic sine transformation

	Tier 1 cumulative imports (asinh)	Tier 2 cumulative imports (asinh)
Panel A: Post 2009		
Minors × Post 2009	-0.0422*	-0.0535**
	(0.022)	(0.027)
Share minors × Post 2009	-0.1705**	-0.3594***
	(0.071)	(0.094)
Percentile share minors × Post 2009	-0.0162***	-0.0269***
	(0.006)	(0.007)
Observations	428,076	428,076
Panel B: Post 2012		
Minors × Post 2012	-0.0445*	-0.0488
	(0.025)	(0.030)
Share minors × Post 2012	-0.1525*	-0.3262***
	(0.089)	(0.115)
Percentile share minors × Post 2012	-0.0163**	-0.0250***
	(0.007)	(0.008)
Observations	353,203	353,203
Firm FE	Yes	Yes
Controls × Year	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report the two cumulative import outcomes from Table B4, replacing $\log(1 + Y)$ with $\operatorname{asinh}(Y)$. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. This table is reported only as a transformation sensitivity for the cumulative outcomes; it does not change their interpretation as stock-like supporting outcomes rather than direct annual adoption effects. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Controls follow Table B4.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B6: Four annual outcomes under the original-paper control set

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Panel A: Post 2009				
Minors × Post 2009	0.0013**	0.0020**	0.0125**	0.0108***
	(0.001)	(0.001)	(0.005)	(0.004)
Share minors × Post 2009	-0.0019	-0.0024	0.0495**	0.0339**
	(0.002)	(0.003)	(0.022)	(0.017)
Percentile share minors × Post 2009	0.0003*	0.0004	0.0052***	0.0036***
	(0.000)	(0.000)	(0.002)	(0.001)
Observations	428,076	428,076	428,076	428,076
Panel B: Post 2012				
Minors × Post 2012	0.0015**	0.0026**	0.0156***	0.0163***
	(0.001)	(0.001)	(0.006)	(0.004)
Share minors × Post 2012	-0.0039*	-0.0036	0.0531**	0.0431**
	(0.002)	(0.003)	(0.023)	(0.019)
Percentile share minors × Post 2012	0.0003	0.0005*	0.0058***	0.0048***
	(0.000)	(0.000)	(0.002)	(0.001)
Observations	353,203	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes	Yes
Baseline educ. × Year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual Tier 1 importer, annual Tier 2 importer, occupation proxy (flexible), and TWFE worker. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the same three treatment codings as Table 16. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Relative to Table 16, this specification retains baseline average years of education, measured in 2008, interacted with year dummies, but omits baseline firm size × year.

Source: Authors' calculations using QP and CI data, 2002–2019.

Table B7: Five annual and cumulative outcomes under the original-paper control set

	Tier 1 importer	Tier 2 importer	ICT worker	Tier 1 cum. imports (log)	Tier 2 cum. imports (log)
Panel A: Post 2009					
Minors × Post 2009	0.0013**	0.0020**	0.0108***	0.0693***	0.1371***
	(0.001)	(0.001)	(0.004)	(0.018)	(0.024)
Share minors × Post 2009	-0.0019	-0.0024	0.0339**	-0.1580**	-0.3320***
	(0.002)	(0.003)	(0.017)	(0.062)	(0.081)
Percentile share minors × Post 2009	0.0003*	0.0004	0.0036***	0.0116**	0.0213***
	(0.000)	(0.000)	(0.001)	(0.005)	(0.006)
Observations	428,076	428,076	428,076	428,076	428,076
Panel B: Post 2012					
Minors × Post 2012	0.0015**	0.0026**	0.0163***	0.0802***	0.1520***
	(0.001)	(0.001)	(0.004)	(0.021)	(0.028)
Share minors × Post 2012	-0.0039*	-0.0036	0.0431**	-0.1668**	-0.3440***
	(0.002)	(0.003)	(0.019)	(0.079)	(0.099)
Percentile share minors × Post 2012	0.0003	0.0005*	0.0048***	0.0146**	0.0254***
	(0.000)	(0.000)	(0.001)	(0.006)	(0.008)
Observations	353,203	353,203	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes	Yes
Region × Year FE	Yes	Yes	Yes	Yes	Yes
Baseline educ. × year	Yes	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual Tier 1 importer, annual Tier 2 importer, TWFE worker, cumulative Tier 1 imports $\log(1+\cdot)$, and cumulative Tier 2 imports $\log(1+\cdot)$. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the same three treatment codings as Table B4. The cumulative columns are best interpreted as stock-like supporting outcomes rather than as direct annual adoption effects. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Relative to Table B4, this specification retains baseline average years of education, measured in 2008, interacted with year dummies, but omits baseline firm size × year.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B8: Binary-exposure estimates under alternative baseline control sets

	Spec 1	Spec 2	Spec 3
Panel A: Post 2009			
Tier 1 importer	-0.0002	-0.0003	-0.0002
	(0.001)	(0.001)	(0.001)
Tier 2 importer	-0.0000	-0.0001	-0.0002
	(0.001)	(0.001)	(0.001)
ICT worker	0.0116***	0.0098***	0.0089**
	(0.004)	(0.004)	(0.004)
Tier 1 cum. imports (log)	-0.0400*	-0.0433**	-0.0453**
	(0.020)	(0.021)	(0.021)
Tier 2 cum. imports (log)	-0.0508**	-0.0530**	-0.0545**
	(0.025)	(0.025)	(0.026)
Observations	428,076	427,835	427,744
Panel B: Post 2012			
Tier 1 importer	-0.0003	-0.0004	-0.0004
	(0.001)	(0.001)	(0.001)
Tier 2 importer	0.0003	0.0002	0.0000
	(0.001)	(0.001)	(0.001)
ICT worker	0.0163***	0.0146***	0.0137***
	(0.004)	(0.004)	(0.004)
Tier 1 cum. imports (log)	-0.0421*	-0.0462*	-0.0483**
	(0.024)	(0.024)	(0.025)
Tier 2 cum. imports (log)	-0.0466	-0.0494*	-0.0525*
	(0.028)	(0.029)	(0.029)
Observations	353,203	353,031	352,944

Notes: Difference-in-differences linear model with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. All entries report the coefficient on the binary exposure interacted with the relevant post indicator for the five outcomes reported in Table B4. Spec 1 includes baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies. Spec 2 adds baseline log firm age × year. Spec 3 adds baseline workforce average age × year to Spec 2. Baseline firm age and baseline workforce average age are measured in 2008. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B9: DiD estimates without sector-year, region-year, or baseline interacted controls

	Robot Dummy	Cumulative Robot imports (log)	ICT Worker
Panel A: Post at 2009			
Minors x Post 2009	0.0014**	0.0697***	0.0123***
	(0.001)	(0.018)	(0.004)
Share minors x Post 2009	-0.0015	-0.1166**	0.0488***
	(0.001)	(0.058)	(0.016)
Percentile share minors x Post 2009	0.0003*	0.0124***	0.0042***
	(0.000)	(0.005)	(0.001)
Observations	435,898	435,898	435,898
Panel B: Post at 2012			
Minors x Post 2012	0.0016**	0.0795***	0.0182***
	(0.001)	(0.021)	(0.004)
Share minors x Post 2012	-0.0029	-0.1167*	0.0559***
	(0.002)	(0.069)	(0.018)
Percentile share minors x Post 2012	0.0003	0.0150***	0.0056***
	(0.000)	(0.006)	(0.001)
Observations	362,557	362,557	362,557
Time FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes

Notes: Difference-in-differences linear model with firm and year fixed effects and firm-clustered standard errors. Columns report annual robot importer, cumulative robot imports $\log(1+\cdot)$, and TWFE worker. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. No sector × year effects, region × year effects, or baseline interacted controls are included.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B10: Probit coefficients for annual robot-importer and TWFE-worker outcomes

	Robot Dummy	ICT Worker
Panel A: Post at 2009		
Minors x Post 2009	0.0403	0.1018***
	(0.119)	(0.039)
Share minors x Post 2009	5.0859	0.0000
	(5.069)	(0.000)
Observations	423,407	428,076
Panel B: Post at 2012		
Minors x Post 2012	0.0833	0.1396***
	(0.096)	(0.040)
Share minors x Post 2012	-1.0709	0.0000
	(2.783)	(0.000)
Observations	349,559	353,203

Notes: Probit models for the two binary annual outcomes, estimated without firm fixed effects and with firm-clustered standard errors. Columns report annual robot importer and TWFE worker. Entries report raw probit coefficients on the interaction between the stated exposure measure and the relevant post indicator. Rows report the binary exposure and the continuous pre-reform share of minors only. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. Controls include year dummies, sector indicators, region indicators, baseline average years of education, and baseline firm size, with baseline variables measured in 2008.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B11: Firm-level DiD estimates on the 2005–2015 window

	Robot Dummy	Cumulative Robot imports (log)	ICT Worker
Panel A: Post at 2009			
Minors x Post 2009	0.0005	-0.0289*	0.0068**
	(0.001)	(0.015)	(0.003)
Share minors x Post 2009	0.0005	-0.0744	0.0318***
	(0.002)	(0.051)	(0.011)
Percentile share minors x Post 2009	0.0001	-0.0097**	0.0025***
	(0.000)	(0.004)	(0.001)
Observations	292,182	292,182	292,182
Panel B: Post at 2012			
Minors x Post 2012	0.0003	-0.0286	0.0124***
	(0.001)	(0.018)	(0.004)
Share minors x Post 2012	-0.0019	-0.0611	0.0415**
	(0.002)	(0.066)	(0.016)
Percentile share minors x Post 2012	0.0001	-0.0093*	0.0036***
	(0.000)	(0.005)	(0.001)
Observations	237,551	237,551	237,551
Time FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated on the 2005–2015 window with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual robot importer, cumulative robot imports $\log(1 + \cdot)$, and TWFE worker. Entries report the coefficient on the interaction between the stated exposure measure and the relevant post indicator. Rows report the binary exposure, the continuous pre-reform share of minors, and the percentile-share treatment. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2005–2015.

Table B12: PPML sensitivity for annual robot importing and cumulative robot-import value

	Robot Importer	Cum. Robot Value (levels)
Panel A: Post 2009		
Minors × Post	0.2578	0.0675
	(0.296)	(0.338)
Panel B: Post 2012		
Minors × Post	0.2436	0.0646
	(0.243)	(0.340)
Estimator	PPML	PPML

Notes: Poisson pseudo-maximum-likelihood models estimated with firm and common year fixed effects and firm-clustered standard errors. Columns report the annual robot-importer indicator and cumulative robot-import value in levels. Entries report the coefficient on the binary exposure interacted with the relevant post indicator. Under the PPML mean specification, $\exp(\beta) - 1$ is the multiplicative change in the conditional mean associated with the interaction term: for the importer column, in the conditional probability of importing; for the cumulative-value column, in expected cumulative robot import value. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. The models also include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies. Unlike the baseline linear specification, these regressions do not absorb sector × year and region × year effects, and the PPML estimator drops firms with separated all-zero outcome paths, so the estimation sample is much smaller. This table is therefore reported only as a narrow functional-form sensitivity for the binary exposure.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B13: Annual comparison across import tiers and occupation proxies

	Tier 1 importer	Tier 2 importer	Tier 3 importer	Occ. strict proxy	Occ. flexible proxy
Panel A: Post 2009					
Minors × Post 2009	-0.0002	-0.0000	0.0005	0.0000	0.0082
	(0.001)	(0.001)	(0.001)	(.)	(0.005)
Observations	428,076	428,076	428,076	202,010	419,226
Panel B: Post 2012					
Minors × Post 2012	-0.0003	0.0003	0.0009	0.0024	0.0127**
	(0.001)	(0.001)	(0.002)	(0.004)	(0.006)
Observations	353,203	353,203	353,203	193,366	346,850
Firm FE	Yes	Yes	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated using the binary exposure only and including firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Entries report the coefficient on the binary exposure interacted with the relevant post indicator. Columns report annual Tier 1 importer, annual Tier 2 importer, annual Tier 3 importer, occupation proxy (strict), and occupation proxy (flexible). The strict occupation proxy is available only on a narrower occupation-harmonized support, which changes the observation counts and leaves the Panel A standard error not estimable in the current sample. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B14: Cumulative comparison across import tiers

	Tier 1 cumulative imports (log)	Tier 2 cumulative imports (log)	Tier 3 cumulative imports (log)
Panel A: Post 2009			
Minors × Post 2009	-0.0400*	-0.0508**	-0.0150
	(0.020)	(0.025)	(0.033)
Observations	428,076	428,076	428,076
Panel B: Post 2012			
Minors × Post 2012	-0.0421*	-0.0466	-0.0235
	(0.024)	(0.028)	(0.037)
Observations	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes

Notes: Difference-in-differences linear model corresponding to Equation 5, estimated using the binary exposure only and including firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Entries report the coefficient on the binary exposure interacted with the relevant post indicator. Columns compare cumulative import outcomes for Tier 1, Tier 2, and Tier 3. These variables are reported because robot acquisitions may be too lumpy to appear clearly in annual importer status, so a reform effect could in principle emerge in accumulated import stocks instead; they should therefore be read as stock-like proxies for accumulated robot-related imports rather than as direct measures of annual adoption events. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B15: Supplementary event-study lead-test p-values for the four main annual adjustment outcomes

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Lead-test p-value (panel A event study)	0.016	0.289	0.560	0.193

Notes: Entries are Panel A event-study lead-test p-values under the binary exposure. Each p-value tests the joint null that all pre-reform exposure × year coefficients are equal to zero in the corresponding event-study regression. The Tier 1 importer and TWFE entries come from the main event-study diagnostics, while the Tier 2 importer and occupation proxy (flexible) entries come from the corresponding annual robot-measure comparison event-study file. All diagnostics use the same fixed-effects structure and baseline education- and baseline size-by-year controls as the main event-study specification.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B16: Threshold sensitivity for binary exposure in four annual adjustment outcomes

	Tier 1 importer	Tier 2 importer	Occ. flexible proxy	ICT worker
Panel A: Post 2009				
Any minors × Post 2009	-0.0002	-0.0000	0.0082	0.0116***
	(0.001)	(0.001)	(0.005)	(0.004)
Share minors ≥ 5% × Post 2009	-0.0005**	-0.0008	0.0017	0.0061*
	(0.000)	(0.001)	(0.005)	(0.003)
Share minors ≥ 10% × Post 2009	0.0000	0.0002	-0.0038	0.0019
	(0.000)	(0.000)	(0.006)	(0.004)
Observations	428,076	428,076	428,076	428,076
Panel B: Post 2012				
Any minors × Post 2012	-0.0003	0.0003	0.0127**	0.0163***
	(0.001)	(0.001)	(0.006)	(0.004)
Share minors ≥ 5% × Post 2012	-0.0008*	-0.0011*	0.0027	0.0096**
	(0.000)	(0.001)	(0.005)	(0.004)
Share minors ≥ 10% × Post 2012	-0.0008	-0.0007	-0.0051	0.0031
	(0.001)	(0.001)	(0.006)	(0.005)
Observations	353,203	353,203	353,203	353,203
Firm FE	Yes	Yes	Yes	Yes
Controls × Year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences linear model with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Columns report annual Tier 1 importer, annual Tier 2 importer, occupation proxy (flexible), and TWFE worker. Entries report the coefficient on the interaction between the stated threshold-based binary exposure and the relevant post indicator. Rows redefine the binary treatment using increasingly demanding pre-reform exposure thresholds: any minors, pre-reform share of minors $\geq 5\%$, and pre-reform share of minors $\geq 10\%$. These codings use the same underlying pre-reform share measure as the continuous treatment in the main table, but they isolate progressively smaller treated groups and should therefore be read as threshold sensitivity rather than as the same estimand as the any-minors dummy in Table 16. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2002–2019.

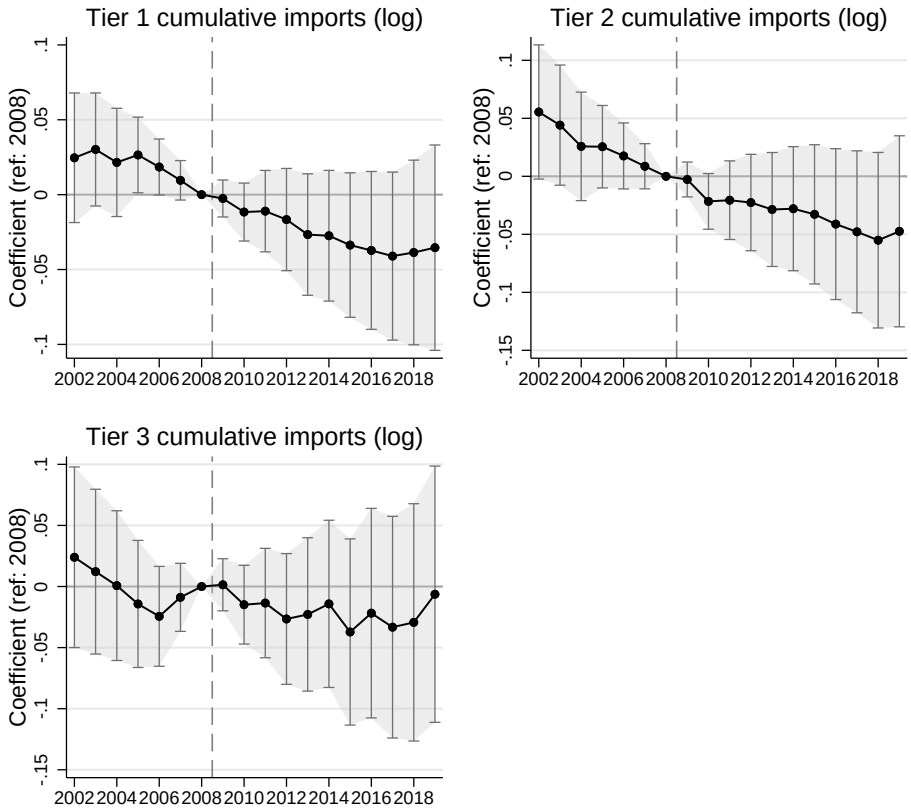


Figure B1: Event study: cumulative import outcomes across Tier 1, Tier 2, and Tier 3 (Panel A)

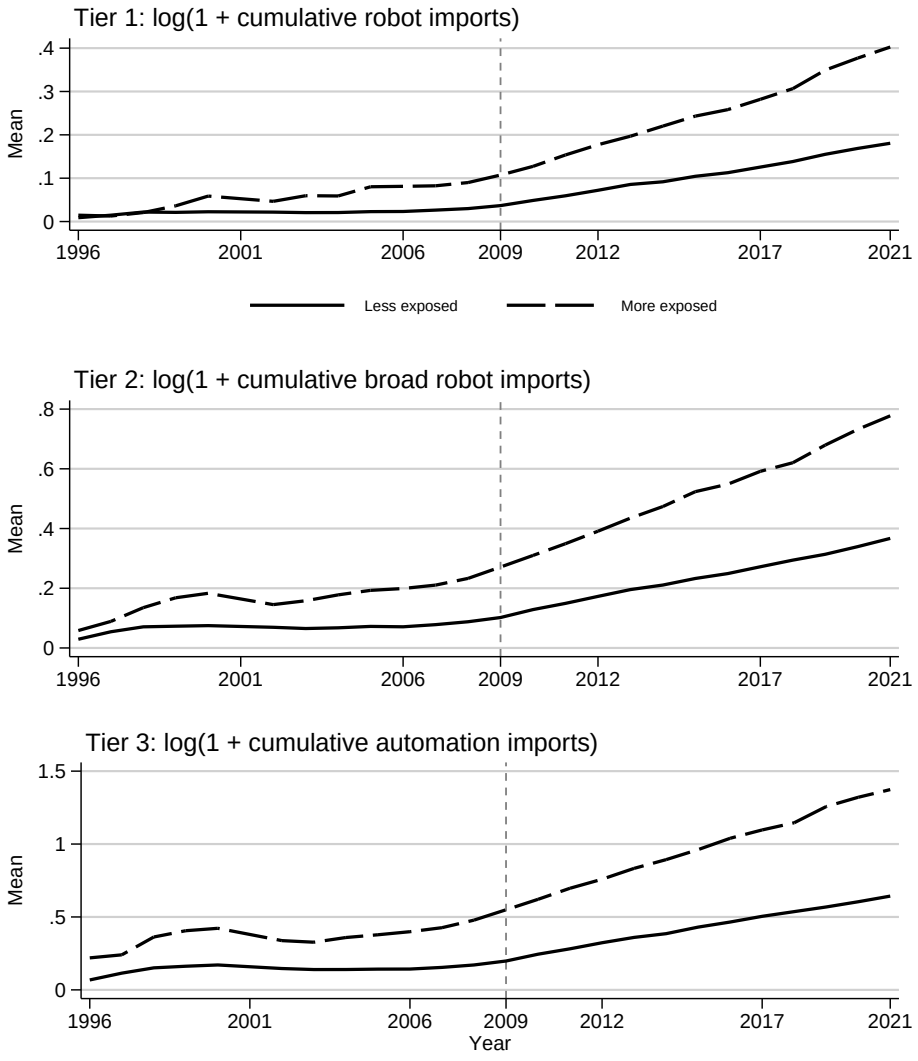


Figure B2: Mean cumulative automation outcomes for more and less exposed firms, 1996–2021 (Tier 1, Tier 2, Tier 3; exposure defined as any worker under age 18 during 2006–2008)

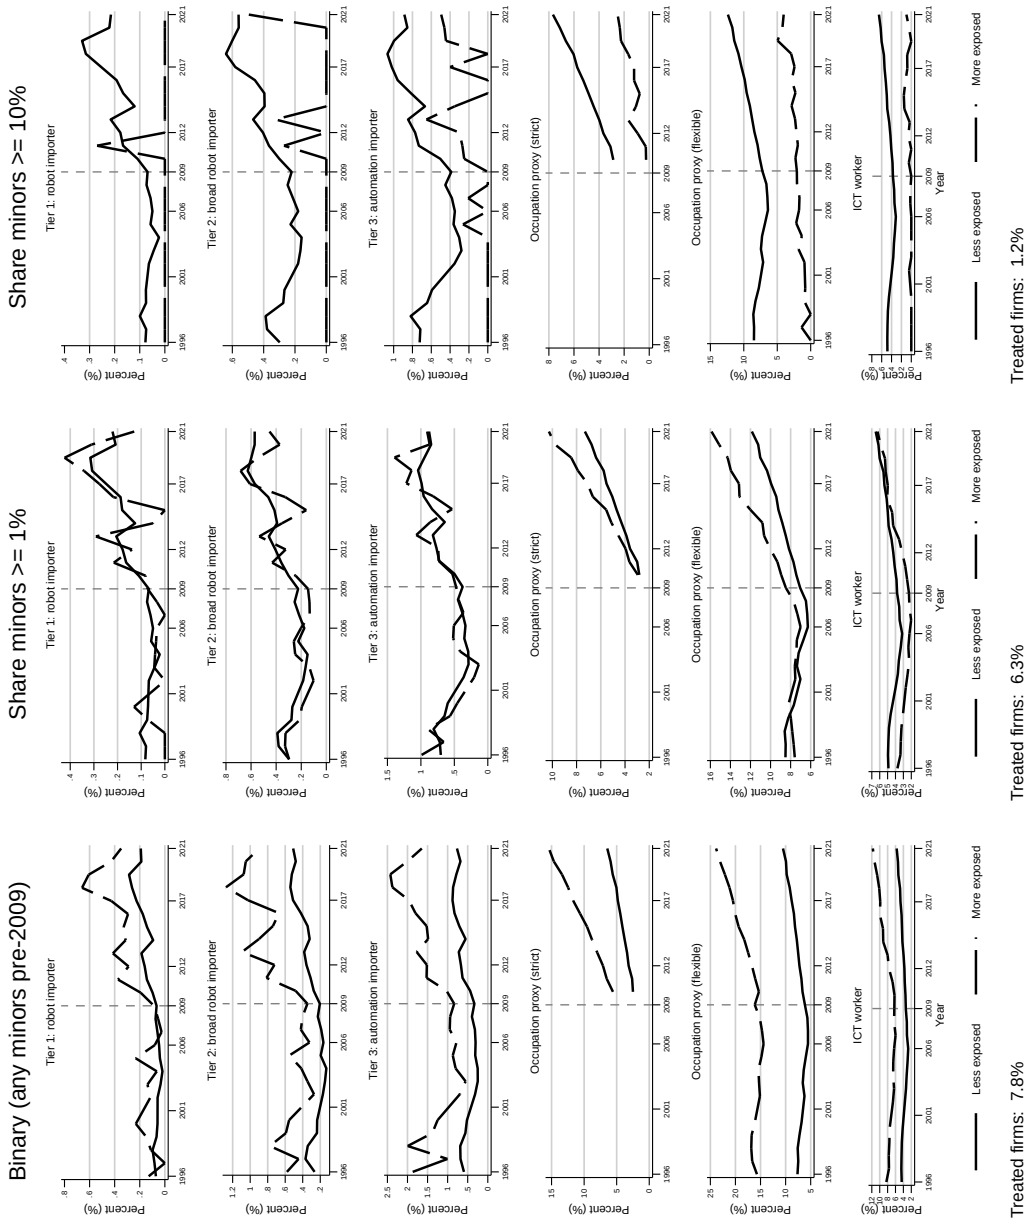


Figure B3: Annual outcomes under alternative exposure definitions (binary, 1%, and 10% cutoffs)

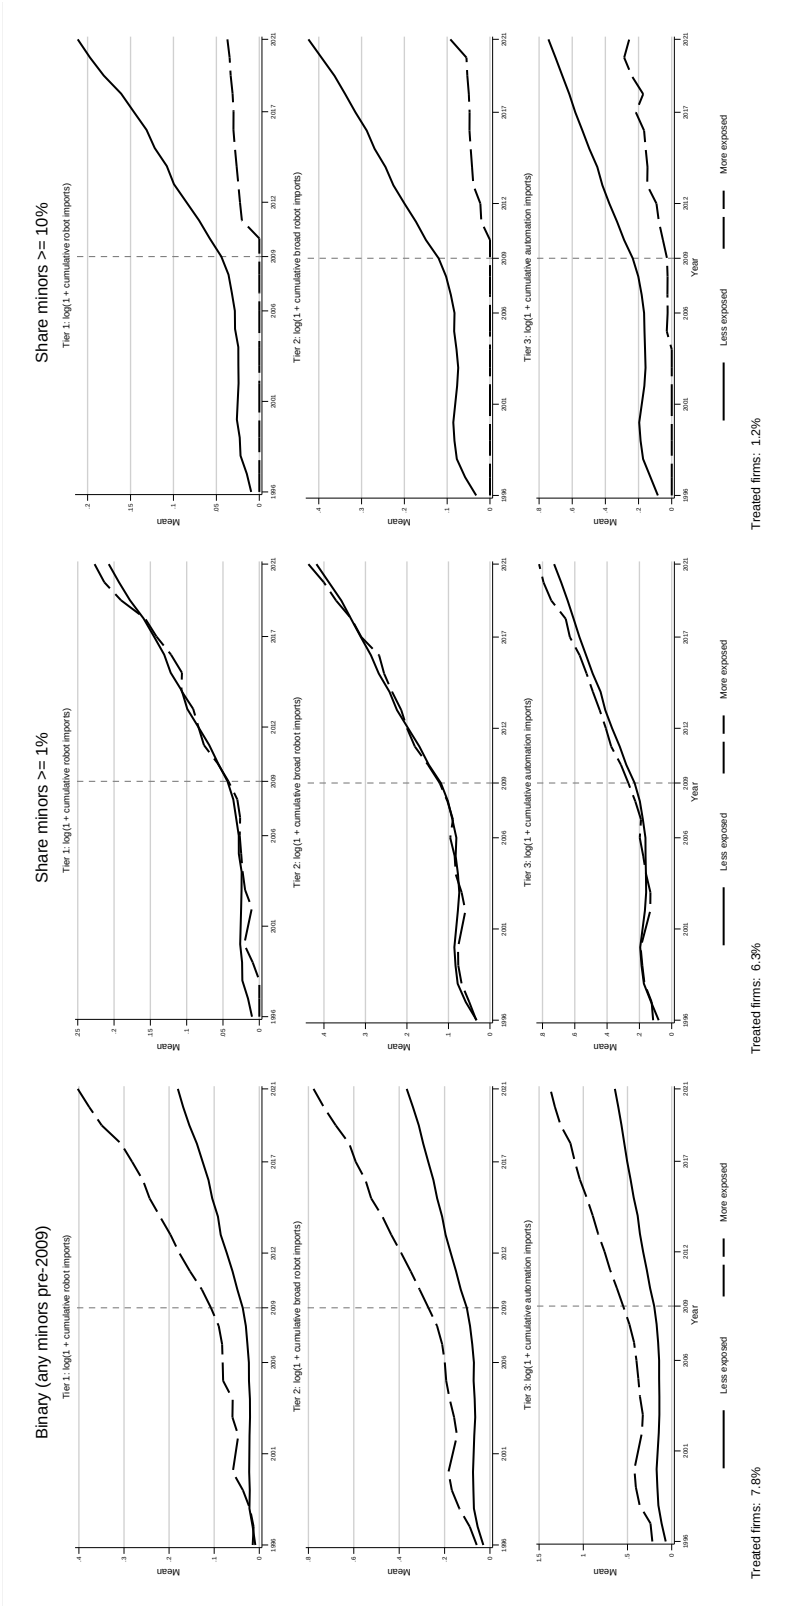


Figure B4: Cumulative outcomes under alternative exposure definitions (binary, 1%, and 10% cutoffs)

Table B17: Binary-exposure estimates under imputation and trend-augmented specifications

	TWFE 2009	Imp. 2009	Trend 2009	TWFE 2012	Imp. 2012	Trend 2012
Robot importer	-0.0002	0.0000	0.0009	-0.0003	-0.0003	0.0004
	(0.0007)	(0.0008)	(0.0009)	(0.0009)	(0.0010)	(0.0012)
Log(1 + robot imports)	-0.0400*	-0.0362	-0.0081	-0.0421*	-0.0399	0.0029
	(0.0205)	(0.0234)	(0.0135)	(0.0235)	(0.0260)	(0.0162)
ICT worker	0.0116***	0.0143***	-0.0031	0.0163***	0.0183***	0.0086**
	(0.0038)	(0.0042)	(0.0033)	(0.0042)	(0.0046)	(0.0035)

Notes: Binary exposure only. Rows report the coefficient on the binary exposure interacted with the relevant post indicator for annual Tier 1 importer, cumulative Tier 1 imports $\log(1 + \cdot)$, and ICT worker. Columns labelled 'TWFE' reproduce the firm fixed-effects difference-in-differences linear model from Table B4. 'Imp.' reports the single-cohort imputation ATT, coding more exposed firms as first treated in 2009 or 2012 and less exposed firms as never treated. 'Trend' augments the firm fixed-effects specification with an exposure-specific linear time trend. Columns labelled 2009 correspond to the Panel A/Post-2009 sample; columns labelled 2012 correspond to the Panel B/Post-2012 sample. All specifications use firm-clustered standard errors and the same baseline education- and baseline size-by-year controls as Table B4. *Source:* Authors' calculations using QP and CI data, 2002–2019.

B.3 Placebo exercises

This appendix reports a broader placebo battery. It distinguishes between placebo *exposure* variables and placebo *timing*. The evidence is mixed. The negative-control exposure placebos are mostly null for the robot outcomes, and the horse-race specifications leave the true binary-exposure coefficients largely unchanged, which is somewhat reassuring against the narrow concern that the robot results are driven simply by generic reliance on nearby-age labour. But those same negative-control exposures remain clearly positive for TWFE, and the fake-timing grid produces non-zero pre-reform coefficients for the annual importer together with smaller positive pseudo-effects for TWFE in some pre-reform years. The placebo evidence should therefore be read as cautionary rather than as a validation exercise.

Table B18: Matched common-support balance diagnostics

	Panel A		Panel B	
Variable	Support	Matched	Support	Matched
Pre-period observations	0.293	0.042	0.261	0.059
Pre ICT worker	0.141	0.013	0.137	0.014
Pre log(1+cum. robots)	0.075	0.030	0.070	0.033
Pre importer rate	0.469	0.000	0.467	0.001
Pre log(1+imports)	0.477	0.002	0.476	0.005
Baseline log employment	1.119	0.067	1.063	0.070
Baseline log sales	0.496	0.002	0.476	0.004
Baseline education	0.067	0.005	0.067	0.020
Firm age in 2008	0.072	0.013	0.077	0.009

Notes: Entries report absolute standardized mean differences for the binary-exposure sample. 'Support' restricts the sample to exact common-support cells defined by baseline sector, region, and size class. 'Matched' reports the same diagnostics after three-nearest-neighbour matching with replacement on pre-2009 outcomes and firm characteristics: the number of pre-period observations, TWFE employment, cumulative robot imports, general import propensity, import value, employment, sales, education, and firm age. Lower values indicate better balance between more-exposed and less-exposed firms.
Source: Authors' calculations using QP and CI data, pre-2009 window.

Table B19: Matched common-support estimates for binary exposure

	Baseline TWFE	Support TWFE	Matched TWFE	Matched TWFE+Trend
Panel A: Post-2009 sample				
Robot importer	-0.0002	-0.0004	0.0002	0.0003
	(0.0007)	(0.0007)	(0.0008)	(0.0010)
Log(1 + robot imports)	-0.0400*	-0.0375*	0.0161	-0.0103
	(0.0205)	(0.0203)	(0.0219)	(0.0155)
ICT worker	0.0116***	0.0117***	0.0023	-0.0029
	(0.0038)	(0.0038)	(0.0049)	(0.0043)
Panel B: Post-2012 sample				
Robot importer	-0.0003	-0.0005	0.0004	-0.0006
	(0.0009)	(0.0009)	(0.0009)	(0.0013)
Log(1 + robot imports)	-0.0421*	-0.0395*	0.0204	-0.0136
	(0.0235)	(0.0233)	(0.0260)	(0.0173)
ICT worker	0.0163***	0.0162***	0.0056	0.0085*
	(0.0042)	(0.0042)	(0.0057)	(0.0048)

Notes: Binary exposure only. Rows report the coefficient on the binary exposure interacted with the relevant post indicator for annual Tier 1 importer, cumulative Tier 1 imports $\log(1 + \cdot)$, and ICT worker. 'Baseline TWFE' reproduces the full-sample firm fixed-effects difference-in-differences model. 'Support TWFE' re-estimates that model after restricting the sample to exact common-support cells defined by baseline sector, region, and size class. 'Matched TWFE' reweights control firms using three-nearest-neighbour matches with replacement on pre-2009 outcomes and firm characteristics. 'Matched TWFE+Trend' adds an exposure-specific linear time trend to the matched specification. Panel A reports the Post-2009 sample; Panel B reports the Post-2012 sample. All columns use firm-clustered standard errors.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table B20: Standard-error comparison under alternative clustering

	ICT Worker	Cum. Robot Imports (log)
Panel A: Post 2009		
$\hat{\beta}$	0.0116**	-0.0400
SE (firm cluster)	0.0038	0.0205
p-value (firm)	0.002	0.051
SE (sector× region)	0.0034	0.0187
p-value (cell)	0.001	0.033
Panel B: Post 2012		
$\hat{\beta}$	0.0163***	-0.0421
SE (firm cluster)	0.0042	0.0235
p-value (firm)	0.000	0.074
SE (sector× region)	0.0041	0.0180
p-value (cell)	0.000	0.020

Notes: This table compares firm-level clustered standard errors (baseline) with sector × NUTS 2 region cell-level clustered standard errors for the binary-exposure TWFE specification. Columns report ICT worker and cumulative robot imports $\log(1 + \cdot)$. Rows report the coefficient estimate, the standard error under each clustering choice, and the corresponding p-value. Panel A includes firms that existed in 2009 and uses the Post-2009 coding. Panel B includes firms that existed in 2012 and uses the Post-2012 coding. *Source:* Authors' calculations using QP and CI data, 2002–2019.

Table B21: Sector–region cell-level DiD estimates

	Tier 1 importer share	Tier 2 importer share	Occ. flexible proxy share	ICT worker share
Panel A: Post 2009				
High exposure × Post 2009	-0.0008***	-0.0014***	0.0031	0.0014
	(0.0002)	(0.0003)	(0.0112)	(0.0026)
Panel B: Post 2012				
High exposure × Post 2012	-0.0005	-0.0009**	-0.0088**	0.0049
	(0.0003)	(0.0004)	(0.0036)	(0.0033)
Observations	770	770	770	770
Cell FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Baseline educ. × year	Yes	Yes	Yes	Yes

Notes: Difference-in-differences model corresponding to Equation 9, estimated with cell and year fixed effects and cell-clustered standard errors. Each cell is a manufacturing sector × NUTS 2 region combination. The four columns mirror the firm-level outcomes in Table 16. "High exposure" is a dummy equal to 1 when the pre-reform share of workers under 18 in the cell exceeds 14% over 2006–2008. Panel A includes firms that existed in 2009; Panel B includes firms that existed in 2012. All specifications include baseline average education in 2008, interacted with year dummies. This table is estimated on 770 cell-year observations and should not be read as a like-for-like replication of Table 16.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table C1: Placebo estimates using negative-control exposure variables

Design	Sample	Robot importer	Log(1 + robot imports)	ICT worker
Female share × Post	Post 2009	0.0010**	0.0685***	0.0018
		(0.0005)	(0.0147)	(0.0028)
Female share × Post	Post 2012	0.0018***	0.0866***	0.0016
		(0.0006)	(0.0184)	(0.0034)
Age 19–21 share × Post	Post 2009	-0.0010	-0.0344	0.0240***
		(0.0008)	(0.0241)	(0.0053)
Age 19–21 share × Post	Post 2012	-0.0014	-0.0470	0.0240***
		(0.0011)	(0.0294)	(0.0060)
Age 20–22 share × Post	Post 2009	-0.0010	-0.0206	0.0234***
		(0.0007)	(0.0214)	(0.0047)
Age 20–22 share × Post	Post 2012	-0.0009	-0.0382	0.0244***
		(0.0009)	(0.0266)	(0.0058)
Horse race: minors × Post	Post 2009	-0.0002	-0.0397*	0.0110***
		(0.0007)	(0.0205)	(0.0038)
Horse race: minors × Post	Post 2012	-0.0003	-0.0413*	0.0157***
		(0.0009)	(0.0236)	(0.0042)
Horse race: age 20–22 × Post	Post 2009	-0.0010	-0.0086	0.0201***
		(0.0007)	(0.0208)	(0.0046)
Horse race: age 20–22 × Post	Post 2012	-0.0009	-0.0250	0.0194***
		(0.0009)	(0.0260)	(0.0057)

Notes: Placebo difference-in-differences linear models corresponding to Equation 5, estimated with firm, sector × year, and region × year fixed effects and firm-clustered standard errors. Entries report the coefficient on the placebo exposure interacted with the relevant post indicator. The table reports four designs: female-share placebo, age 19–21 placebo, age 20–22 placebo, and a horse-race specification that includes both the true binary exposure and the age 20–22 placebo exposure. Panel A includes firms that existed in 2009 and uses the Post-2009 coding; Panel B includes firms that existed in 2012 and uses the Post-2012 coding. All specifications include baseline average years of education and baseline firm size, both measured in 2008, interacted with year dummies.
Source: Authors' calculations using QP and CI data, 2002–2019.

Table C2: Placebo timing estimates with fake reform years in the pre-2009 window

Fake reform year	Robot importer	Log(1 + robot imports)	ICT worker
2004	-0.0010	-0.0070	0.0034
	(0.0008)	(0.0127)	(0.0034)
2005	-0.0014**	-0.0060	0.0047
	(0.0007)	(0.0140)	(0.0030)
2006	-0.0023***	-0.0118	0.0046*
	(0.0007)	(0.0127)	(0.0027)
2007	-0.0024***	-0.0164	0.0030
	(0.0008)	(0.0112)	(0.0025)
2008	-0.0020**	-0.0178*	0.0042
	(0.0008)	(0.0106)	(0.0026)

Notes: Placebo timing exercise restricted to the pre-2009 window. Each row reports the coefficient on the binary exposure interacted with a fake post indicator in a firm fixed-effects difference-in-differences linear model that assigns a placebo reform year between 2004 and 2008. The purpose is to assess whether the pre-reform window already contains differential dynamics between more-exposed and less-exposed firms.
Source: Authors' calculations using QP and CI data, 2002–2008.

C

Minimum wage, robot adoption, and worker adjustment: Evidence from Portugal

C.1 Additional descriptive statistics

Table C3: Descriptive statistics adopters and non-adopters by size

	Non-adopters	Adopters
< 10 workers		
PRCI	0.27	0.04
RTI	0.25	-0.10
PBOE	0.40	0.36
Labour productivity (ths)	23.25	113.27
Number of employees	4.43	5.68
Exporter	0.16	0.44
Tangible assets (ths)	99.40	1073.34
Firm age	14.72	10.45
Worker age	41.87	39.44
Sh. of workers with higher education	0.07	0.35
>= 10 & < 20 workers		
PRCI	0.23	0.05
RTI	0.18	0.03
PBOE	0.40	0.41
Labour productivity (ths)	27.90	62.91
Number of employees	13.53	14.30
Exporter	0.36	0.62
Tangible assets (ths)	356.17	1615.80
Firm age	18.99	14.01
Worker age	41.04	38.84

	Non-adopters	Adopters
Sh. of workers with higher education	0.07	0.21
>= 20 & < 50 workers		
PRCI	0.19	0.08
RTI	0.13	0.08
PBOE	0.41	0.43
Labour productivity (ths)	34.67	51.12
Number of employees	30.51	33.03
Exporter	0.53	0.75
Tangible assets (ths)	1100.95	2059.81
Firm age	23.47	20.57
Worker age	41.17	39.26
Sh. of workers with higher education	0.09	0.18
>= 50 & < 250 workers		
PRCI	0.11	0.06
RTI	0.08	0.04
PBOE	0.42	0.43
Labour productivity (ths)	41.41	56.93
Number of employees	97.53	119.82
Exporter	0.73	0.84
Tangible assets (ths)	4172.23	7377.30
Firm age	29.27	28.89
Worker age	40.97	40.21
Sh. of workers with higher education	0.11	0.16
>= 250 workers		
PRCI	0.06	0.03
RTI	0.02	0.01
PBOE	0.42	0.43
Labour productivity (ths)	58.16	67.57
Number of employees	498.54	753.80
Exporter	0.81	0.92
Tangible assets (ths)	32453.55	37071.45
Firm age	37.68	30.67
Worker age	40.88	39.76
Sh. of workers with higher education	0.14	0.17

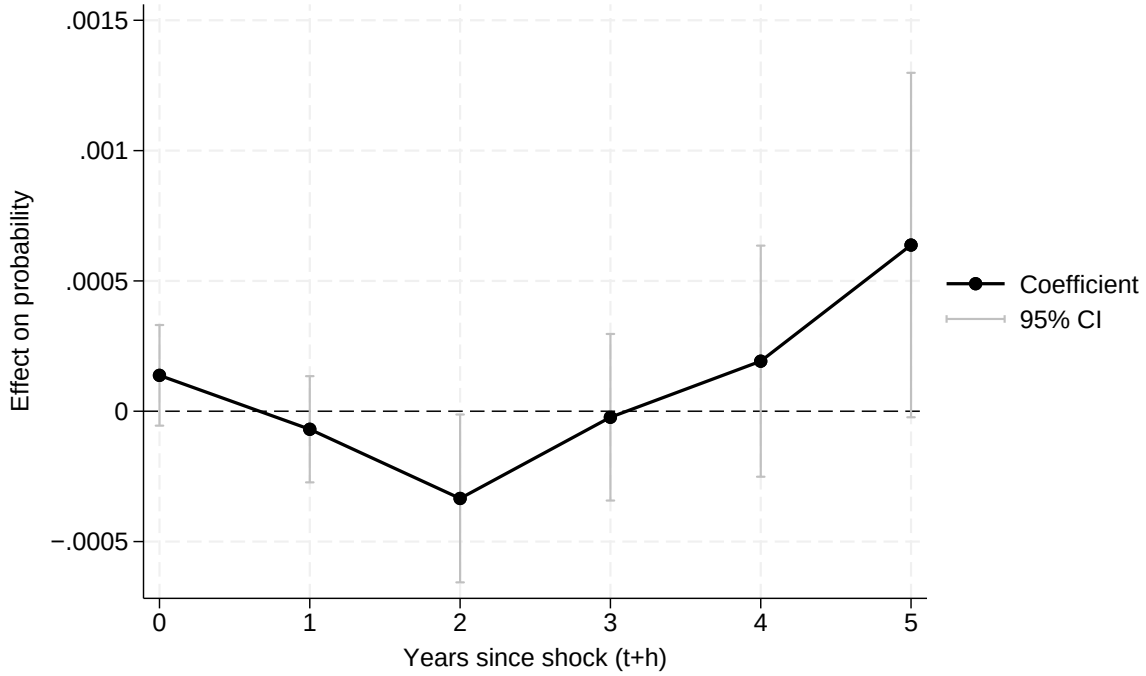
Notes: The table presents sample averages for firm and worker characteristics, split by robot adoption status. "Non-adopters" refers to firms that never imported industrial robots during the sample period, while "Adopters" refers to firms with a positive robot stock. Variables denoted with "(ths)" are expressed in thousands of euros. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Table C4: Within-firm variation in occupational exposure measures (2010)

	Standard Deviation			Multi-Occupation Share	
Variable	Overall	Between	Within	Firms	Employment
RTI	0.636	0.485	0.564	73.6%	95.9%
PBOE	0.101	0.087	0.082	73.6%	95.9%

Notes: This table reports the variance decomposition for the two occupational exposure measures used in the worker-level analysis. To avoid contamination from post-adoption occupational sorting, all statistics are calculated strictly on the 2010 base year cross-section. 'Overall', 'Between', and 'Within' standard deviations are calculated at the firm level. 'Multi-Occupation Share' reports the percentage of firms and the percentage of total workers located in firms that employed workers in more than one 2-digit ISCO-08 occupational category in 2010. RTI stands for Routine Task Index. PBOE stands for Patent-Based Occupational Exposure. *Source:* Authors' calculations based on QP and SCIE.

C.2 Additional results – minimum wage and robot adoption



Notes: The figure plots the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in $t+h$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). See section 4.4 for further details on the specification. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. See Table 23 for further details. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C1: Minimum wage and robot adoption , purchase in t – local projection coefficients

Table C5: Minimum wage and robot adoption with a stable sample

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI_{it}^{pred}$	-0.0017*	-0.0017*	-0.0012	0.0007	0.0025**	0.0041***
	(0.0010)	(0.0009)	(0.0008)	(0.0007)	(0.0010)	(0.0014)
Log(Worker age)	-0.0016	0.0015	0.0023	-0.0014	-0.0014	0.0033
	(0.0034)	(0.0031)	(0.0033)	(0.0040)	(0.0036)	(0.0032)
Log(productivity)	0.0010	0.0001	-0.0007	0.0003	0.0001	0.0000
	(0.0007)	(0.0008)	(0.0009)	(0.0009)	(0.0007)	(0.0008)
% bachelor	0.0031	0.0055	0.0106**	0.0126**	0.0127**	-0.0008
	(0.0041)	(0.0049)	(0.0051)	(0.0055)	(0.0057)	(0.0037)
Wage premium	-0.0033	-0.0012	0.0031	0.0023	0.0025	0.0015
	(0.0023)	(0.0022)	(0.0024)	(0.0026)	(0.0024)	(0.0021)
Observations	78386	78386	78386	78386	78386	78386

Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections considering a stable sample across all periods. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t + h$ to $t + h + 2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). All specifications include firm fixed effects, industry-year fixed effects, and time-varying firm controls (log of average worker age, log labour productivity, and the share of workers with higher education). Standard errors clustered at the firm level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Table C6: Minimum wage and robot adoption with lagged exposure in levels

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI^{pred}_{f,t}$	-0.0000	-0.0001	-0.0002	0.0004	0.0011***	0.0017***
	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0004)	(0.0007)
$PRCI^{pred}_{f,t-1}$	0.0003**	0.0003*	0.0005**	0.0007***	0.0009***	0.0005
	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0003)	(0.0004)
Log(Worker age)	-0.0041**	-0.0014	0.0017	0.0022	0.0003	0.0027
	(0.0020)	(0.0022)	(0.0023)	(0.0026)	(0.0025)	(0.0028)
Log(productivity)	0.0004	0.0000	-0.0005	0.0003	-0.0000	0.0001
	(0.0004)	(0.0005)	(0.0006)	(0.0006)	(0.0006)	(0.0007)
% bachelor	0.0031	0.0036	0.0045	0.0065*	0.0079**	0.0005
	(0.0029)	(0.0034)	(0.0036)	(0.0037)	(0.0037)	(0.0036)
Observations	172092	150317	132200	114930	98434	82565

Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t+h$ to $t+h+2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI^{pred}_{ft}$). As robustness check we include $PRCI^{pred}_{f,t-1}$ to assess whether the results are driven by the incremental shock rather than by the cumulative level of exposure. All specifications include firm fixed effects, industry-year fixed effects, and time-varying firm controls (log of average worker age, log labour productivity, and the share of workers with higher education). Standard errors clustered at the firm level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Table C7: Minimum wage and robot adoption two-way clustered SEs (firm and year)

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI_{it}^{pred}$	0.0000	-0.0000	-0.0001	0.0006	0.0013**	0.0019*
	(0.0002)	(0.0002)	(0.0002)	(0.0003)	(0.0005)	(0.0007)
Log(Worker age)	-0.0040*	-0.0013	0.0018	0.0023	0.0004	0.0027
	(0.0019)	(0.0023)	(0.0021)	(0.0025)	(0.0028)	(0.0022)
Log(productivity)	0.0005	0.0000	-0.0005	0.0003	0.0000	0.0001
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0006)	(0.0007)
% bachelor	0.0031	0.0035	0.0044	0.0065	0.0078	0.0005
	(0.0029)	(0.0038)	(0.0045)	(0.0043)	(0.0046)	(0.0060)
Observations	172092	150317	132200	114930	98434	82565

Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t+h$ to $t+h+2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). All specifications include firm fixed effects, industry-year fixed effects, and time-varying firm controls (log of average worker age, log labour productivity, and the share of workers with higher education). Standard errors clustered at the firm and year level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Table C8: Minimum wage and robot adoption without control variables

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI_{it}^{pred}$	-0.0000	-0.0001	-0.0001	0.0005*	0.0011**	0.0018***
	(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0005)	(0.0007)
Observations	177586	153980	135034	117236	100386	84229

Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t+h$ to $t+h+2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). All specifications include firm fixed effects and industry-year fixed effects. Standard errors clustered at the firm level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Table C9: Minimum wage and robot adoption with lagged control variables

	(1) t	(2) t+1	(3) t+2	(4) t+3	(5) t+4	(6) t+5
$\Delta PRCI_{it}^{pred}$	0.0000	-0.0001	-0.0001	0.0006*	0.0012**	0.0017***
	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0005)	(0.0007)
$Log(Workerage)_{it-1}$	-0.0023	0.0005	0.0012	-0.0007	0.0003	-0.0007
	(0.0019)	(0.0021)	(0.0023)	(0.0023)	(0.0023)	(0.0035)
$Log(productivity)_{it-1}$	0.0003	-0.0001	0.0003	0.0000	-0.0002	0.0000
	(0.0004)	(0.0005)	(0.0006)	(0.0005)	(0.0006)	(0.0006)
$\%bachelor_{it-1}$	0.0029	0.0049	0.0048	0.0042	0.0001	-0.0059*
	(0.0031)	(0.0035)	(0.0036)	(0.0034)	(0.0030)	(0.0034)
Observations	173285	150757	132428	115024	98535	82661

Notes: The table reports the estimated dynamic response (β_h) of robot adoption to minimum wage shocks using local projections. Each column represents a different horizon h. The dependent variable is an indicator equal to 1 if the firm purchases industrial robots in the three-year window from $t + h$ to $t + h + 2$. The primary explanatory variable is the first-difference of the predetermined Potential Relative Cost Increase ($\Delta PRCI_{ft}^{pred}$). All specifications include firm fixed effects, industry-year fixed effects, and time-varying firm controls (log of average worker age, log labour productivity, and the share of workers with higher education) lagged by one year. Standard errors clustered at the firm level are reported in parentheses. Significance levels: * 10%, ** 5%, *** 1%. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

C.3 Descriptive firm-level trajectories around robot adoption

This appendix presents descriptive evidence on how key firm-level outcomes evolve around the adoption of industrial robots. Because robot adoption is an endogenous strategic decision, the patterns documented here are not interpreted as causal effects of automation. Instead, they serve to characterise the type of firms that adopt robots, to illustrate the presence of positive selection into automation, and to provide descriptive context for the worker-level analysis in the main text.

Figures C2–C6 plot descriptive event-study trajectories for a set of core firm outcomes, including labour productivity, value added, employment, labour share, and export intensity, centred on the year of first robot adoption. Across several outcomes, future adopters exhibit stronger trajectories already in the pre-adoption period. Productivity, output, and employment are higher and tend to grow more rapidly before the adoption event, and exporters are disproportionately represented among adopting firms. These patterns indicate positive selection into automation, consistent with the view that firms self-select into robot adoption from an already favourable position.

This positive selection is central for interpreting the firm-level evidence in the paper. Because adopters are on more expansionary trajectories prior to automation, correlations between robot adoption and subsequent firm performance cannot be interpreted as causal effects of automation on productivity, scale, or exporting. For this reason, the main analysis deliberately refrains from making causal claims about firm-level outcomes. Instead, causal identification is achieved at the worker level by exploiting within-firm variation in pre-adoption task exposure combined with firm-year fixed effects, as described in Sections 4.4 and 4.5.2. The descriptive patterns in this appendix therefore serve as selection diagnostics and background context rather than as part of the paper's causal evidence.

The expansionary profile of adopting firms also provides context for interpreting the worker-level results. The increase in job separations among highly exposed workers, particularly those on fixed-term contracts, is not obviously a by-product of broad firm decline. Rather, the descriptive patterns in this appendix suggest that worker displacement occurs around adoption episodes in firms that are, on average, otherwise expanding. This is consistent with within-firm restructuring around automation rather than simple contraction of distressed firms.

The descriptive evidence also helps clarify how to read the Step 1 results. Step 1 estimates how predetermined exposure to minimum-wage-induced labour cost shocks affects the probability of robot adoption over time (Section 4.5.1). The patterns reported here show that adoption is more common among firms that already display stronger pre-adoption trajectories. This suggests that the estimated response to minimum wage exposure is likely to operate through firms with greater scale, productivity, or organisational capacity to absorb the fixed costs of automation. This interpretation should remain suggestive, however, because the firm-level evidence in this appendix is descriptive rather than causal.

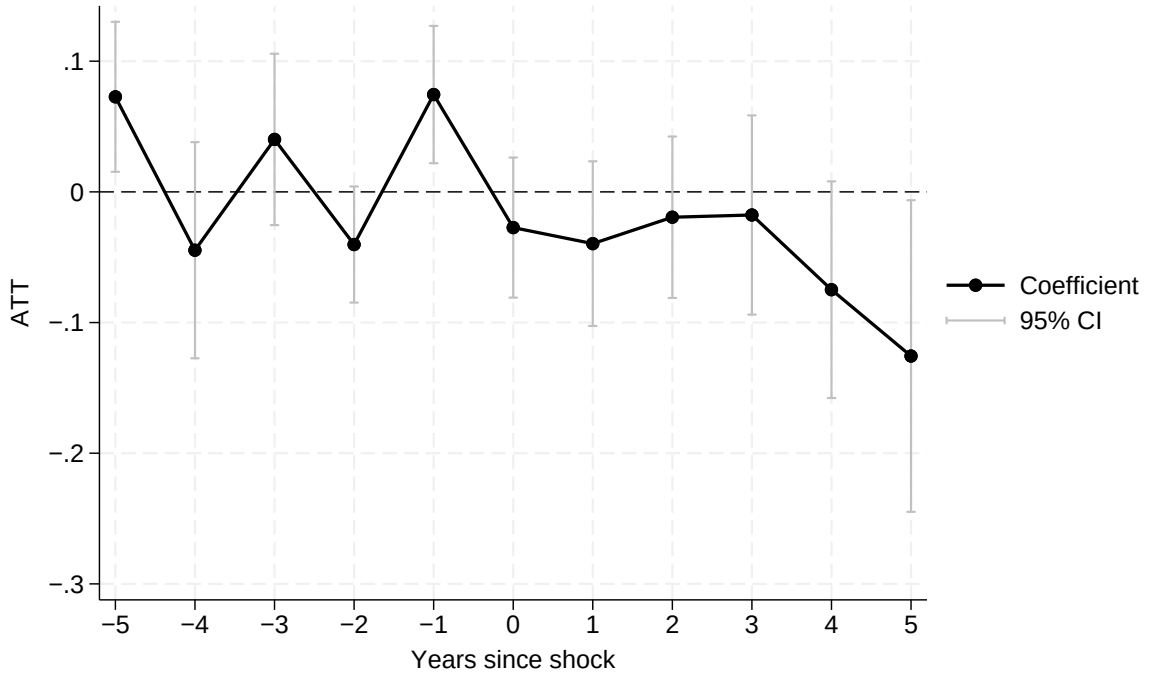
Table C10 complements the event-study figures by reporting descriptive fixed-effects associations between lagged robot adoption and firm outcomes. Consistent with the graphical evidence, the coefficients point to stronger performance among adopters on several dimensions, but these estimates should be

read only as descriptive correlations. Together with the event-study figures, they reinforce the conclusion that robot adoption is preceded by favourable firm dynamics, underscoring the importance of separating selection into automation from the causal effects of automation.

Table C10: Robot adoption and firm-level outcomes (descriptive)

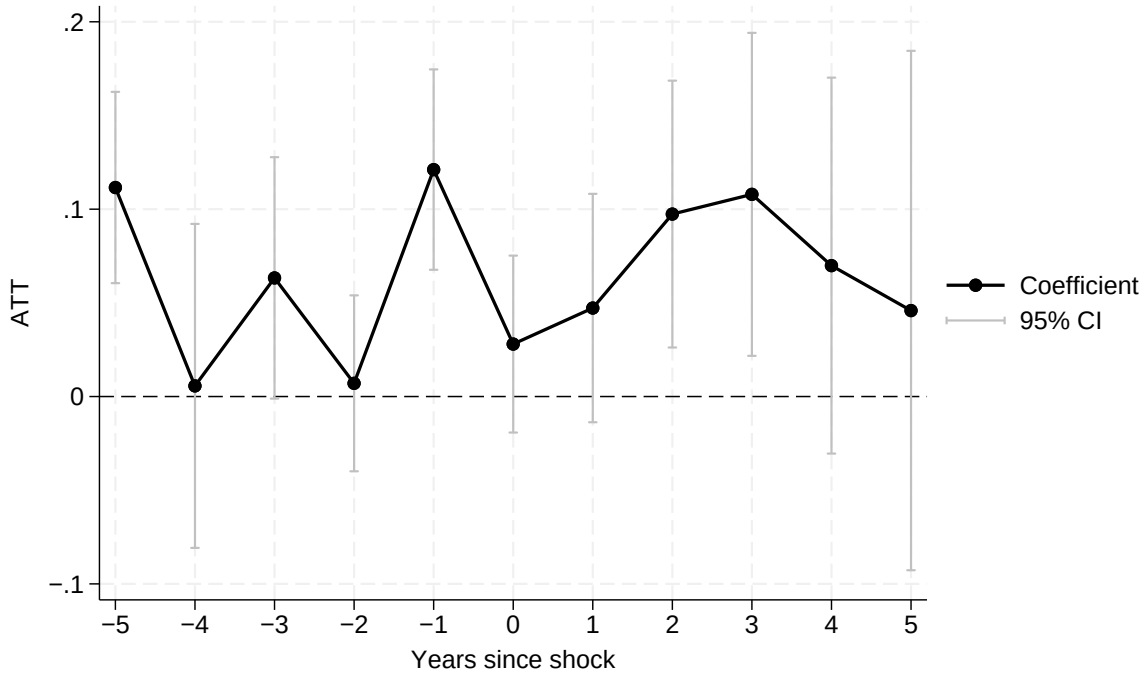
	(1) Productivity	(2) Value added	(3) Employment	(4) Labour share	(5) Exports
$Robots_{it-3}$	0.0520***	0.0766***	0.0770***	-0.0062	0.0271**
	(0.0176)	(0.0194)	(0.0159)	(0.0052)	(0.0124)
$Log(Workerage)_{it}$	-0.0257	-0.2614***	-0.3891***	0.0021	-0.0125
	(0.0181)	(0.0223)	(0.0197)	(0.0059)	(0.0098)
$\%bachelor_{it}$	0.0101	-0.0759***	-0.1597***	0.0172**	0.0079
	(0.0251)	(0.0293)	(0.0273)	(0.0073)	(0.0116)
Observations	202218	202337	207255	189261	208254
R^2	0.719	0.935	0.956	0.663	0.761

Notes: The dependent variables labour productivity, value added, and employment (number of employees) are expressed in logarithms. Labour share is measured in percentage points and the exporter outcome is a dummy variable. Values of labour share below 0 and above 1 are discarded. *L3.Robots* is a binary indicator equal to one if the firm purchased industrial robots three years earlier. All specifications include firm and year fixed effects, as well as fixed effects for firm size class and firm age class. All models control for the log of average worker age and the share of workers with a bachelor's degree. Standard errors are clustered at the firm level and reported in parentheses. *Significance levels:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).



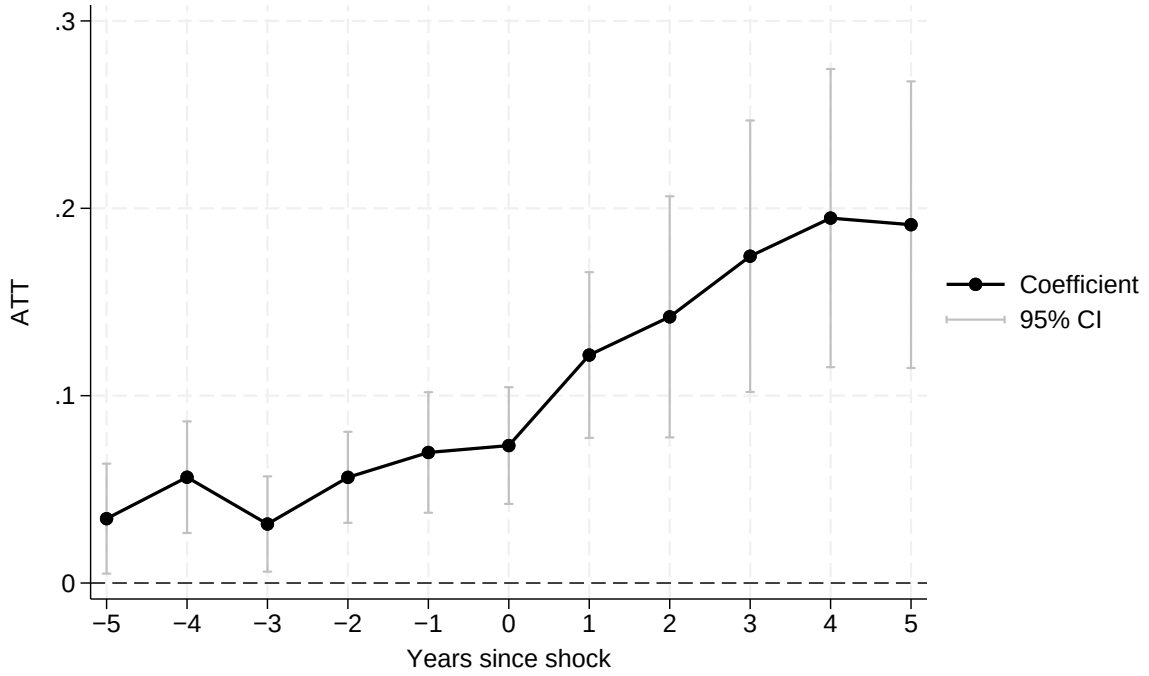
Notes: The dependent variable is labour productivity (in logarithms). The treatment timing variable (*treat1*) identifies the first year a firm purchased robots; firms that never purchased robots are coded as zero, and observations after 2021 are excluded from the analysis. The specifications include firm fixed effects to control for time-invariant unobserved heterogeneity across firms. The models also control for the log of average worker age and the share of workers with a bachelor's degree. The figures report event-study estimates based on the estimator of Callaway and Sant'Anna, 2021. In this appendix, they are presented as descriptive trajectory plots to assess selection into adoption rather than as causal effect estimates. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C2: Robot adoption and productivity – descriptive event-study trajectories



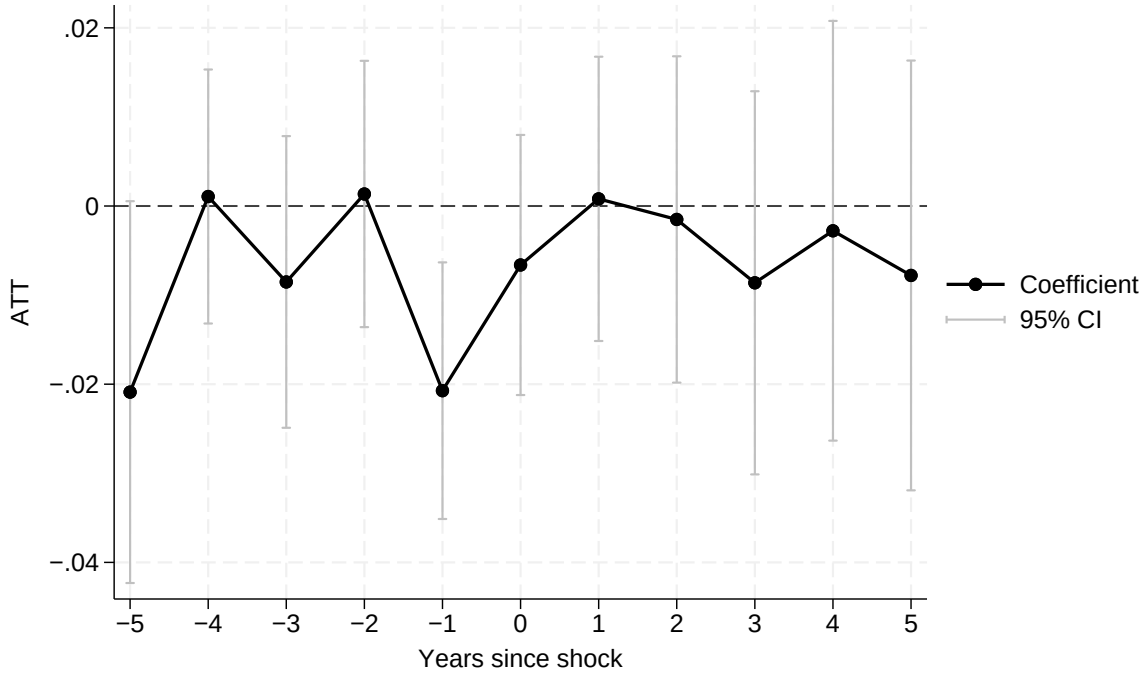
Notes: The dependent variable is value added (in logarithms). The treatment timing variable (*treat1*) identifies the first year a firm purchased robots; firms that never purchased robots are coded as zero, and observations after 2021 are excluded from the analysis. The specifications include firm fixed effects to control for time-invariant unobserved heterogeneity across firms. The models also control for log of average worker age and the share of workers with a bachelor's degree. The figures report event-study estimates based on the estimator of Callaway and Sant'Anna, 2021. In this appendix, they are presented as descriptive trajectory plots to assess selection into adoption rather than as causal effect estimates. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C3: Robot adoption and value added – descriptive event-study trajectories



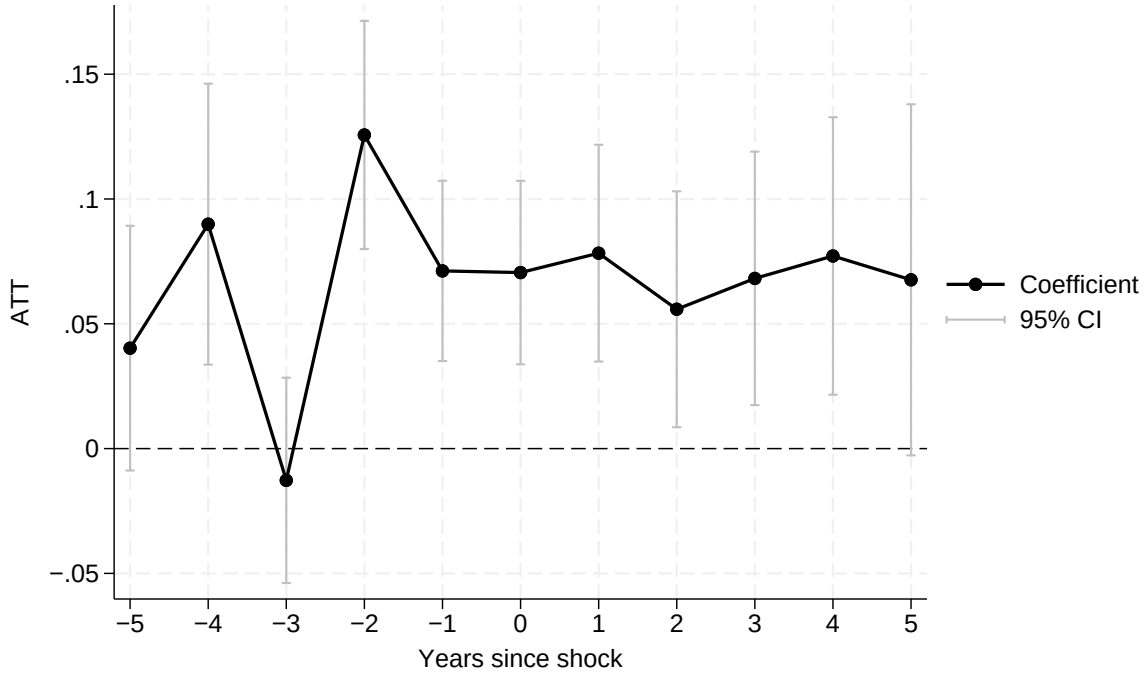
Notes: The dependent variable is employment, measured as the logarithm of the number of workers. The treatment timing variable (*treat1*) identifies the first year a firm purchased robots; firms that never purchased robots are coded as zero, and observations after 2021 are excluded from the analysis. The specifications include firm fixed effects to control for time-invariant unobserved heterogeneity across firms. The models also control for the log of average worker age and the share of workers with a bachelor's degree. The figures report event-study estimates based on the estimator of Callaway and Sant'Anna, 2021. In this appendix, they are presented as descriptive trajectory plots to assess selection into adoption rather than as causal effect estimates. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C4: Robot adoption and employment – descriptive event-study trajectories



Notes: The dependent variable is the labour share, defined as total wages paid divided by value added. The treatment timing variable (*treat1*) identifies the first year a firm purchased robots; firms that never purchased robots are coded as zero, and observations after 2021 are excluded from the analysis. The specifications include firm fixed effects to control for time-invariant unobserved heterogeneity across firms. The models also control for the log of average worker age and the share of workers with a bachelor's degree. The figures report event-study estimates based on the estimator of Callaway and Sant'Anna, 2021. In this appendix, they are presented as descriptive trajectory plots to assess selection into adoption rather than as causal effect estimates. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

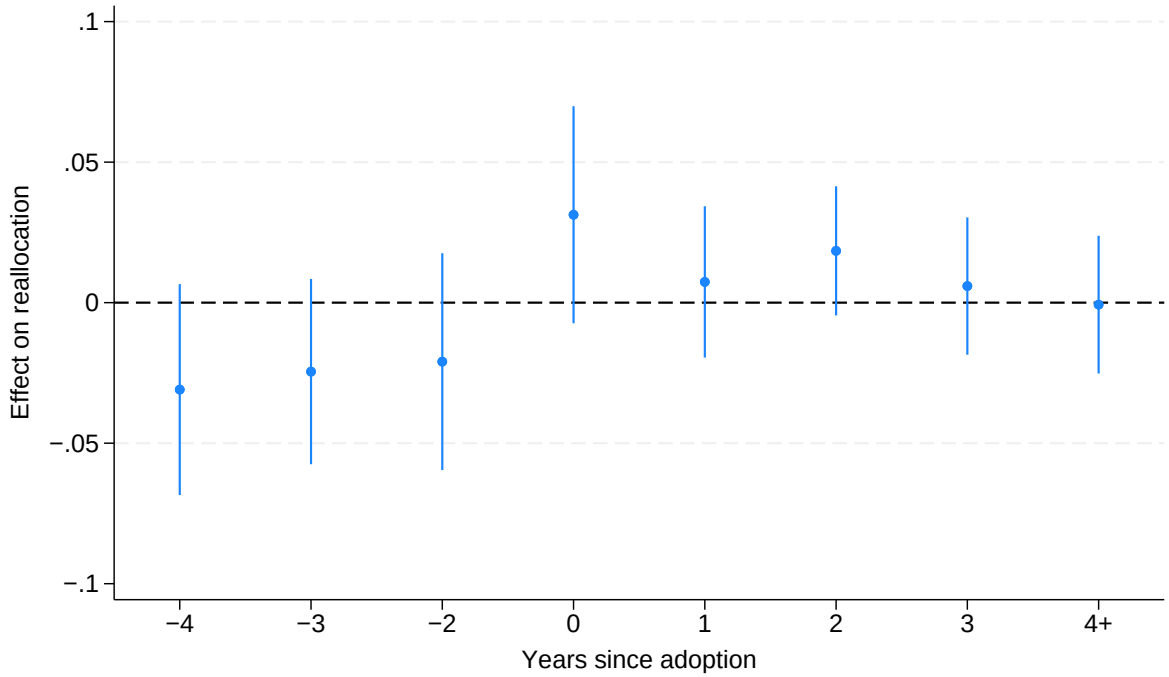
Figure C5: Robot adoption and labour share – descriptive event-study trajectories



Notes: The dependent variable is *exporter*, a binary variable equal to 1 if at least 5% of total sales are to countries other than Portugal, and 0 otherwise. The treatment timing variable (*treat1*) identifies the first year a firm purchased robots; firms that never purchased robots are coded as zero, and observations after 2021 are excluded from the analysis. The specifications include firm fixed effects to control for time-invariant unobserved heterogeneity across firms. The models also control for the log of average worker age and the share of workers with a bachelor's degree. The figures report event-study estimates based on the estimator of Callaway and Sant'Anna, 2021. In this appendix, they are presented as descriptive trajectory plots to assess selection into adoption rather than as causal effect estimates. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C6: Robot adoption and exports – descriptive event-study trajectories

C.4 Additional results – worker level outcomes



Notes: The figure plots the event-study coefficients representing the dynamic reallocation effects of robot adoption, differentiated by workers' pre-adoption automation exposure. Only fixed-term workers are considered. The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Event time zero ($t = 0$) is defined as the year of the firm's first observed robot purchase. See section 3.4. for further details on specifications and sample. Error bars represent 95% confidence intervals based on standard errors clustered at the firm level. *Source:* Authors' calculations based on CI, QP, and SCIE (INE).

Figure C7: Robot adoption, exposure, and worker reallocation – fixed-term workers, event study

Table C11: Robot adoption, exposure (PBOE), and worker wages

	(1)	(2)	(3)	(4)
Robots=1	0.5379***	0.0543**	0.0543**	
	(0.0805)	(0.0268)	(0.0268)	
PBOE	-0.9047***	-0.1237***	-0.1237***	-0.1443***
	(0.0331)	(0.0105)	(0.0105)	(0.0111)
Robots=1 × PBOE	-1.1617***	-0.1086*	-0.1086*	-0.1409***
	(0.1581)	(0.0567)	(0.0567)	(0.0474)
Log(Worker age)	0.5532***	1.0560***	1.0560***	0.9873***
	(0.0101)	(0.0242)	(0.0242)	(0.0271)
Log(Educ)	0.4178***	0.0109***	0.0109***	0.0114***
	(0.0085)	(0.0020)	(0.0020)	(0.0016)
Observations	3164488	3151859	3151859	3099607
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is the log of hourly wages. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C12: Robot adoption, exposure (PBOE), and worker displacement – fixed-term workers

	(1)	(2)	(3)	(4)
Robots=1	0.0431***	-0.0521	-0.0680**	
	(0.0134)	(0.0318)	(0.0304)	
PBOE	0.0709***	-0.0118	0.0007	0.0111
	(0.0095)	(0.0275)	(0.0252)	(0.0264)
Robots=1 × PBOE	-0.1206***	0.0921	0.1176*	0.1452*
	(0.0298)	(0.0686)	(0.0644)	(0.0783)
Log(Worker age)	-0.0223***	0.1219***	0.2407***	0.2364***
	(0.0022)	(0.0269)	(0.0296)	(0.0296)
Log(Educ)	0.0150***	-0.0214***	-0.0096	-0.0105
	(0.0020)	(0.0066)	(0.0062)	(0.0064)
Observations	705906	601092	599510	560714
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Only workers under a fixed-term contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C13: Robot adoption, exposure (PBOE), and worker displacement – all workers

	(1)	(2)	(3)	(4)
Robots=1	0.0217***	-0.0143**	-0.0290***	
	(0.0048)	(0.0068)	(0.0047)	
PBOE	0.0449***	0.0145	0.0336***	0.0436***
	(0.0035)	(0.0100)	(0.0098)	(0.0069)
Robots=1 × PBOE	-0.0591***	0.0205	0.0353***	0.0317**
	(0.0101)	(0.0159)	(0.0099)	(0.0125)
Log(Worker age)	-0.0362***	0.0748***	0.1936***	0.1648***
	(0.0010)	(0.0087)	(0.0095)	(0.0085)
Log(Educ)	0.0103***	-0.0205***	-0.0075***	-0.0090***
	(0.0006)	(0.0022)	(0.0019)	(0.0017)
Observations	3333196	3221097	3219617	3176998
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker separates from the firm in year t+1. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C14: Robot adoption, exposure (PBOE), and worker reallocation – fixed-term workers

	(1)	(2)	(3)	(4)
Robots=1	-0.0378*	0.1088***	0.0946***	
	(0.0229)	(0.0281)	(0.0329)	
PBOE	0.1052***	1.7630***	2.4262***	2.3817***
	(0.0085)	(0.0732)	(0.0962)	(0.0840)
Robots=1 × PBOE	0.0867	-0.2447***	-0.2121***	-0.3762***
	(0.0561)	(0.0660)	(0.0781)	(0.0818)
Log(Worker age)	-0.0025*	-0.0059	0.0011	0.0193
	(0.0013)	(0.0210)	(0.0237)	(0.0183)
Log(Educ)	0.0170***	0.0026	0.0026	0.0053
	(0.0011)	(0.0044)	(0.0054)	(0.0046)
Observations	659942	556927	555822	518391
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Only workers under a fixed-term contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C15: Robot adoption, exposure (PBOE), and worker reallocation – all workers

	(1)	(2)	(3)	(4)
Robots=1	-0.0048	0.0052	0.0025	
	(0.0107)	(0.0107)	(0.0111)	
PBOE	0.0729***	1.3772***	1.5966***	1.5348***
	(0.0053)	(0.0441)	(0.0501)	(0.0458)
Robots=1 × PBOE	0.0134	-0.0130	-0.0089	-0.0227
	(0.0269)	(0.0263)	(0.0275)	(0.0271)
Log(Worker age)	-0.0050***	0.0089	0.0073	0.0126
	(0.0011)	(0.0101)	(0.0104)	(0.0079)
Log(Educ)	0.0137***	0.0028	0.0012	0.0021
	(0.0010)	(0.0018)	(0.0019)	(0.0014)
Observations	3231694	3120141	3119025	3077724
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to an occupation with lower automation exposure in t+1. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.
Source: Authors' calculations based on CI, QP, and SCIE.

Table C16: Robot adoption, exposure, and worker reallocation – any direction – fixed-term

	(1)	(2)	(3)	(4)
Robots=1	0.0069	0.0022	0.0003	
	(0.0110)	(0.0085)	(0.0128)	
RTI	-0.0019	0.0230**	0.0315**	0.0472***
	(0.0019)	(0.0113)	(0.0142)	(0.0151)
Robots=1 × RTI	0.0107	0.0129	0.0057	-0.0028
	(0.0107)	(0.0128)	(0.0141)	(0.0108)
Log(Worker age)	-0.0226***	-0.0652*	-0.0474	-0.0197
	(0.0033)	(0.0393)	(0.0445)	(0.0274)
Log(Educ)	0.0156***	-0.0072	-0.0116	-0.0042
	(0.0028)	(0.0072)	(0.0087)	(0.0073)
Observations	659942	556927	555822	518391
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to a different occupation in t+1. Only workers under a fixed-term contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C17: Robot adoption, exposure, and worker reallocation – any direction – all workers

	(1)	(2)	(3)	(4)
Robots=1	0.0045	0.0069	0.0059	
	(0.0063)	(0.0069)	(0.0083)	
RTI	-0.0054***	0.0104	0.0102	0.0227***
	(0.0008)	(0.0079)	(0.0084)	(0.0048)
Robots=1 × RTI	0.0076*	0.0080	0.0070	0.0042
	(0.0040)	(0.0081)	(0.0085)	(0.0064)
Log(Worker age)	-0.0332***	-0.0457*	-0.0517**	-0.0360**
	(0.0027)	(0.0247)	(0.0256)	(0.0182)
Log(Educ)	0.0119***	-0.0067*	-0.0093**	-0.0026
	(0.0020)	(0.0037)	(0.0043)	(0.0027)
Observations	3231694	3120141	3119025	3077724
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to a different occupation in t+1. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.

Table C18: Robot adoption, exposure, and worker reallocation – any direction – open-ended

	(1)	(2)	(3)	(4)
Robots=1	0.0044	0.0090	0.0076	
	(0.0059)	(0.0079)	(0.0089)	
RTI	-0.0061***	0.0068	0.0056	0.0201***
	(0.0007)	(0.0099)	(0.0105)	(0.0046)
Robots=1 × RTI	0.0063*	0.0045	0.0040	0.0027
	(0.0033)	(0.0087)	(0.0091)	(0.0073)
Log(Worker age)	-0.0262***	-0.0012	-0.0035	0.0043
	(0.0024)	(0.0230)	(0.0236)	(0.0167)
Log(Educ)	0.0112***	-0.0109**	-0.0117**	-0.0028
	(0.0021)	(0.0048)	(0.0054)	(0.0037)
Observations	2571752	2489841	2489003	2444735
Worker FE	No	Yes	Yes	Yes
Firm FE	No	No	Yes	Yes
Firm x Year FE	No	No	No	Yes

Notes: The dependent variable is an indicator equal to 1 if the worker remains employed at the same firm but moves to a different occupation in t+1. Only workers under an open-ended contract are considered. Standard errors clustered at the firm level in parentheses. Significance levels: * 10%, ** 5%, *** 1%.

Source: Authors' calculations based on CI, QP, and SCIE.