ECONSTOR
Make Your Publications Visible.

A Service of

ZBW Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Pulito, Giuseppe; Pytlikova, Mariola; Schroeder, Sarah; Lodefalk, Magnus

Working Paper

Who Adopts AI? Evidence on Firms, Technologies and Workers

GLO Discussion Paper, No. 1732

Provided in Cooperation with:
Global Labor Organization (GLO)

Suggested Citation: Pulito, Giuseppe; Pytlikova, Mariola; Schroeder, Sarah; Lodefalk, Magnus (2026) : Who Adopts AI? Evidence on Firms, Technologies and Workers, GLO Discussion Paper, No. 1732, Global Labor Organization (GLO), Essen

This Version is available at:
https://hdl.handle.net/10419/339715

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Who Adopts AI? Evidence on Firms, Technologies and Workers*

Giuseppe Pulito[1] Mariola Pytlikova[2] Sarah Schroeder[3] Magnus Lodefalk[4†]

March 27, 2026

Using two waves of nationally representative Danish firm surveys linked to employer–employee administrative registers, we study how adoption varies across artificial intelligence (AI) and related advanced technologies. We show that AI adoption is highly technology-specific. While firm size and digital infrastructure predict adoption broadly, workforce composition operates through distinct channels: STEM-educated workforces predict core AI adoption, whereas non-STEM university-educated workforces are associated with generative AI adoption, indicating different human capital complementarities. The factors associated with adoption differ from those predicting deployment breadth: firm size and digital maturity matter for both, whereas workforce composition primarily predicts adoption alone. Machine learning and natural language processing are deployed across multiple business functions, whereas other advanced technologies remain concentrated in specific operational domains. Individual-level evidence provides a foundation for these patterns, with awareness of workplace AI usage concentrated among managers and high-skilled workers. Self-reported AI knowledge is higher among younger and more educated individuals. Finally, commonly used occupational AI exposure measures vary substantially in their ability to predict observed adoption, with benchmark-based measures outperforming patent-based and LLM-focused alternatives. These findings show that treating AI as a monolithic category obscures economically meaningful variation in who adopts, what they deploy, and how well existing measures capture it.

Keywords: Artificial Intelligence; Technology Adoption; Digitalisation; Human capital; AI Exposure Measures.
JEL Codes: D24, J23, J62, O33

*We thank Hildegunn Kyvik Nordas for comments on the survey design and the manuscript, and participants at AIAS, the Colloquium on Personnel Economics (Trondheim, 2025), CReAM/RFBerlin, AI-Econ Lab, the AIAS/PIREAU Conference on "AI and Inequality" and CBS/DST Business Statistics Conference 2025 for helpful suggestions. We also thank Radovan Kubovic, Madalina Suman and Mikkel Braas Skytte for their research assistance. Pulito acknowledges financial support from the ROCKWOOL Foundation Berlin (Project No. 2002), Pytlikova from "New Technologies and Changes in Education, Research and Labor Market" (CZ.02.01.01/00/23_025/0008693, co-funded by the EU), Schroeder from PIREAU and the Carlsberg Foundation, and Lodefalk from the Torsten Söderberg Foundation (grants E46/21, ET3/23) and WASP-HS (grant 805). The usual disclaimer applies.

†[1]ROCKWOOL Foundation Berlin, gp@rfberlin.com. [2]CERGE-EI, Charles University and the Economics Institute of the Czech Academy of Sciences, and Aarhus Institute of Advanced Studies (AIAS) at Aarhus University, mariola.pytlikova@cerge-ei.cz. [3]Aarhus University and Ratio Institute, sschroeder@econ.au.dk. [4]Orebro University, Ratio, Global Labor Organization, magnus.lodefalk@oru.se.

1 Introduction

Artificial intelligence (AI) capabilities are advancing rapidly, and firm-level adoption is accelerating across economies. However, measuring adoption precisely remains challenging, as reported rates differ widely depending on survey design, sample composition, and how AI is defined. In the United States, fewer than 6 per cent of firms reported using any AI-related technology in the 2018 Annual Business Survey (McElheran et al. 2024), a figure that had risen only modestly by 2024, with the Census Bureau's real-time tracking estimating firm-level AI use at around 5–6 per cent (Bonney et al. 2024). In Europe, the Eurostat enterprise survey documents an increase from 8 to 20 per cent between 2023 and 2025, with Denmark leading at 42 per cent.[1] In a recent multi-country survey of senior business executives, Yotzov et al. 2026 report that 69 per cent of firms across the US, UK, Germany, and Australia currently use some AI technology, though over 90 per cent of these firms report no material impact on employment or productivity to date. Cross-country evidence from eleven OECD economies confirms that adoption varies markedly by firm size and sector, with larger and more digitally mature firms consistently at the frontier (Calvino and Fontanelli 2023). These differences reflect not only variation in diffusion across countries and firms, but also differences in how AI is conceptualised and measured. This raises a question the literature has only begun to address: which AI technologies do firms actually adopt, and do the organisational capabilities supporting adoption differ?

Why differences in the technologies firms adopt, and in the organisational capabilities required to implement them, matter is well established in the economics of technology diffusion. New technologies deliver productivity gains only when firms make complementary investments in management practices, workforce skills, and business-process redesign (Bresnahan et al. 2002; Brynjolfsson et al. 2021), and recent evidence from US manufacturing confirms that this logic extends to AI, with early adopters experiencing a J-curve in productivity (McElheran et al. 2025). If the required complements differ across technologies, then treating AI as a single binary variable risks conflating distinct diffusion processes. The rapid spread of generative AI makes this concern concrete: while generative AI is often available through low-cost, publicly accessible tools and requires minimal firm-level infrastructure, core AI technologies such as machine learning or computer vision typically demand deliberate investment in data systems, technical expertise, and organisational redesign. As a result, different technologies are likely to diffuse along distinct margins across firms and what predicts adoption of one technology class may not predict the adoption of another.

We address this question using two waves of nationally representative firm surveys (2023, 2024) and complementary worker surveys, conducted by Statistics Denmark and

1. Other Scandinavian countries, such as Sweden and Finland, also ranked among the top adopters in 2025 (Eurostat 2026).

linked to rich employer–employee administrative registers. The surveys capture not only whether firms adopt AI, but also distinguish between different technology classes, including core AI technologies (machine learning, natural language processing, machine vision, voice recognition), generative AI, and a set of complementary advanced technologies, and the business functions in which the technology is used. The two survey waves span the generative AI transition: by 2024, roughly half of surveyed firms report using generative AI, allowing us to examine this emerging technology class alongside more established ones.

A key feature of our analysis is that firm characteristics and workforce composition variables from administrative registers are measured in 2022, prior to the adoption observed in 2023 and 2024. This predetermined covariate structure strengthens the interpretation of the register-based associations, though two predictors, the digitalisation index and cloud importance, are drawn from the survey and are therefore contemporaneous with adoption. We exploit this structure to examine AI adoption along several dimensions. First, we study the extensive margin (whether firms adopt AI at all) and show that the organisational capabilities associated with adoption differ across core AI, generative AI, and other advanced technologies. Second, we analyse the intensive margin among adopters, in terms of the breadth of technologies deployed, to assess whether the factors that predict initial adoption also predict deeper technological engagement. Third, we examine how AI diffuses across business functions within firms, distinguishing technologies that spread broadly from those that remain tied to specific operational domains. Fourth, we turn to individual-level survey data, which serve two purposes: they benchmark firm-reported adoption against workers' own workplace experience and reveal substantial heterogeneity in AI use and knowledge across occupations, education levels, and demographic groups that is invisible in firm-level data alone. Finally, by combining survey-based adoption measures with detailed occupational composition from administrative registers, we assess how well widely used AI exposure indices predict observed firm adoption.

Our central finding is that AI adoption patterns are technology-specific. The organisational capabilities associated with adoption vary across technology types and adoption margins, implying that treating AI as a monolithic category masks economically meaningful differences in diffusion patterns. Five findings support this interpretation. First, core AI technologies and generative AI have distinct predictor profiles. While firm size and digital infrastructure predict adoption across technology types, workforce composition operates through different channels: a STEM-university-educated workforce is associated with core AI adoption, whereas a non-STEM university-educated workforce predicts generative AI adoption, suggesting that the two technology classes complement different types of human capital. Second, machine learning and natural language processing diffuse across multiple business functions within adopting firms, consistent with the behaviour expected of a general-purpose technology; other advanced technologies remain tightly coupled to specific operational domains. Third, the factors associated with initial adoption

differ from those predicting deployment breadth among adopters: firm size and digital maturity matter for both, whereas workforce composition – particularly STEM education – is primarily linked to the adoption decision alone, indicating that crossing the adoption threshold and scaling AI within the firm draw on distinct organisational capabilities. Fourth, individual-level evidence complements the firm-level analysis. Worker-reported AI use closely aligns with firm-reported adoption and is concentrated among managers and professionals, supporting the use of managerial respondents in firm surveys. Patterns in self-assessed AI knowledge provide a micro-level foundation for the firm-level results: knowledge is higher among younger and more educated individuals, including both STEM and non-STEM university graduates, and varies systematically across demographic groups. Together, these patterns indicate that workforce composition—especially the availability of educated and digitally capable workers—is a key correlate of AI adoption and its use within firms.. Fifth, commonly used AI exposure indices vary substantially in their ability to predict observed firm adoption, with benchmark-based capability measures outperforming patent-based and large-language-model-focused alternatives, underscoring that how AI exposure is measured can shape what researchers find.

The paper makes three contributions. First, we show that the firm characteristics associated with AI adoption are technology-specific, using a classification that distinguishes core AI technologies (machine learning, NLP, machine vision, voice recognition), generative AI, and other advanced technologies. Recent US evidence documents that adoption of industrial AI correlates with firm size, digital infrastructure, and firm strategies (McElheran et al. 2024), that early adopters experience a productivity J-curve shaped by management practices and firm age (McElheran et al. 2025), and European work identifies similar adoption patterns at a broader level (Rammer et al. 2022). These studies treat AI as a single category, reflecting the data available at the time (McElheran et al. 2024; McElheran et al. 2025). We build on this work by decomposing AI into distinct technology types, showing that workforce composition operates through different channels for each: STEM-university-educated workforces are associated with core AI adoption, while non-STEM university-educated workforces predict generative AI adoption. Our two survey waves span the generative AI transition, making these data among the first to compare established and emergent AI adoption within the same analytical framework. Our analysis complements Humlum and Vestergaard 2024, who study individual-level ChatGPT adoption, by examining firm-level decisions across the full spectrum of AI technologies, and extends recent evidence on how different uses of AI shape labour demand (Aghion et al. 2025) by identifying what firm characteristics are associated with deployment of these technologies in the first place.

Second, we document that specific AI technologies, machine learning and natural language processing, are associated with AI use across multiple business functions within adopting firms, consistent with the broader applicability theorised for general-purpose

technologies (Brynjolfsson et al. 2021). Other advanced technologies are associated with use in specific functional domains only, suggesting that the degree of cross-functional reach varies substantially within the broader family of AI-related technologies.

Third, we offer one of the first systematic validations of occupational AI exposure indices against observed firm adoption in a linked employer–employee setting. Despite relying on the same underlying occupational structure, these measures differ substantially in their association with observed adoption, with indices grounded in expert assessments of AI capabilities (Felten et al. 2021; Engberg et al. 2024) outperforming patent-based (Webb 2019) and LLM-focused (Eloundou et al. 2023) alternatives. These results have direct implications for the growing empirical literature that uses exposure measures to study AI's effects on workers and local labour markets.

The remainder of the paper proceeds as follows. Section 2.1 describes the survey and administrative data. Section 2.2 presents descriptive evidence on AI adoption. Sections 3 and 4 examine firm-level and worker-level correlates of AI adoption, respectively. Section 5 validates commonly used AI exposure measures. Finally, Section 6 concludes.

2 Data and Descriptive Statistics

2.1 Data Sources

2.1.1 The Danish AI & Digitisation Survey

Our analysis draws on two waves of the Danish firm-level survey on Digitisation and AI, conducted by Statistics Denmark. The questionnaire collects detailed information on firms' awareness, use, and attitudes towards artificial intelligence in the workplace. The survey was administered in mixed mode, combining online questionnaires (CAWI) with telephone interviews (CATI), with fieldwork and invitations managed entirely by Statistics Denmark. Respondents were instructed to answer on behalf of their company rather than in a personal capacity and were shown the standard Statistics Denmark confidentiality and data-protection statement. Further details on the questionnaire design and implementation are provided in Appendix Section A.

The survey targets private-sector firms with at least five employees and is based on a stratified sampling design by 18 industry groups and five firm-size classes, ensuring representativeness along both dimensions. The statistical unit is the legal entity (CVR level), and the population is defined based on the Business Statistics Register. The first wave, fielded in summer 2023 and yielded 2,326 completed responses, corresponding to a response rate of 23.1%. The second wave was fielded in autumn 2024 using the same sampling frame and survey design and yielded 1,907 completed responses, corresponding to a response rate of 18.9%. Statistics Denmark provides survey weights that correct for differential non-response and ensure representativeness of the target population.

We additionally draw on two waves of a representative individual-level survey on artificial intelligence conducted by Statistics Denmark. The survey was administered as a web-based questionnaire (CAWI), with fieldwork and invitations managed entirely by Statistics Denmark. It targets the Danish resident population aged 18–64 and is based on a simple random sample drawn from the population register, ensuring representativeness of the working-age population. The first wave was fielded in 2023, during which 8,814 individuals were contacted and 1,725 completed responses were collected, corresponding to a response rate of 19.6%. The second wave was fielded in 2024 using the same sampling frame and survey design; a comparable number of individuals were contacted, yielding 1,648 completed responses. For both waves, Statistics Denmark provides survey weights that correct for differential non-response and align the sample with population characteristics.

The questionnaire focuses exclusively on the use of AI in the workplace. It contains detailed information on workers' awareness of AI, AI use in job-related tasks, individual's AI knowledge and attitudes towards the use of AI at work, including perceived implications for job content and skill requirements.

2.1.2 Register Data

A key advantage of our setting is the ability to link survey responses with Statistics Denmark's administrative registers at both the firm and worker level. Firm registers provide economy-wide coverage of both the private and public sectors and contain information on industry affiliation, ownership, accounts, exports, balance sheets, revenues, capital stock, and value added. On the labour market side, the Integrated Database for Labour Market Research (IDA) offers universal linked employer–employee data (LEED) since the early 1980s, covering individual demographics, education, detailed occupation codes, wages, and employment histories. Taken together, these resources allow us to study AI adoption with exceptional granularity, connecting firm-level technology choices to rich workforce composition and outcomes.

To ensure population representativeness, Statistics Denmark provides survey weights based on the sampling stratification that adjust for differential response probabilities. To address selection into survey participation beyond the stratification variables and to flexibly account for observable determinants of response, we leverage rich administrative register data to estimate firms' response probabilities and construct inverse probability weights that correct for survey non-response, as described in Appendix C. Table 1 reports summary statistics for the survey respondents, presented in both unweighted and weighted form, alongside the corresponding population means drawn from administrative registers. The comparison shows that the weighting procedure brings the survey sample much closer to the population distribution, particularly with respect to firm size, revenues, and trade activities. This highlights both the importance of applying weights in the analysis and

Table 1: Summary Statistics for the Firm Sample

	Respondents (n=3,998)		Respondents (weighted)		Population (n=43,605)	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**
No. of employees	90.47	370.93	38.19	223.38	36.27	214.10
Revenue (M)	375.40	5,671.80	146.03	2,920.04	135.49	3,759.99
Value added (M)	95.40	507.34	35.95	266.14	40.36	1,291.11
Physical capital stock (M)	23.76	306.80	9.58	165.05	13.89	792.41
Intangible capital stock (M)	0.07	1.26	0.03	0.66	0.05	1.99
Firm age	24.02	19.54	18.58	15.97	19.70	17.31
Exports (M)	170.69	4,090.12	56.41	2,086.28	65.09	3,398.58
Imports (M)	263.68	8,850.30	82.83	4,635.39	66.30	3,284.95
Foreign ownership	0.16	0.37	0.11	0.31	0.10	0.29
Digitalisation index	12.45	4.53	11.71	4.85	12.45	4.53
Cloud importance (indicator)	0.80	0.40	0.76	0.43	0.80	0.40
Yearly wage (th.)	403.87	172.67	369.83	169.61	362.49	272.53
Hourly wage	259.82	83.48	243.93	80.41	242.71	139.35
Female workers	0.36	0.25	0.36	0.28	0.38	0.29
Immigrant workers	0.13	0.18	0.14	0.21	0.13	0.21
Average age (workers)	42.48	6.92	41.51	8.05	41.36	8.13
Younger than 36	0.36	0.21	0.39	0.25	0.39	0.25
Between 36 and 54	0.40	0.15	0.38	0.18	0.38	0.19
Older than 55	0.24	0.16	0.23	0.19	0.23	0.19
No secondary education	0.17	0.16	0.19	0.19	0.19	0.18
Bachelor's or higher	0.36	0.28	0.34	0.30	0.33	0.29
Master's or higher	0.13	0.18	0.12	0.19	0.11	0.18
STEM educated	0.25	0.24	0.24	0.25	0.22	0.25
University STEM education	0.09	0.16	0.08	0.16	0.07	0.14
Union membership	0.63	0.20	0.58	0.23	0.58	0.23
Managers	0.07	0.08	0.06	0.09	0.06	0.10
Professionals	0.35	0.33	0.33	0.34	0.32	0.34
Clerical/Administrative	0.12	0.17	0.12	0.17	0.11	0.17
Sales/Services	0.11	0.22	0.13	0.24	0.17	0.27
Skilled trades	0.14	0.27	0.15	0.28	0.16	0.30
Production workers	0.20	0.28	0.20	0.29	0.17	0.26

Notes: This table reports descriptive statistics for the firm-level data. "Respondents" are firms that completed the AI survey in 2023–2024. Columns labeled "Respondents (weighted)" apply inverse-probability weights to correct for survey nonresponse, as described in Appendix C. The "Population" comprises all non-agricultural firms in Denmark with at least five employees, mirroring the sample restrictions used in the survey. Monetary values are nominal; "M" denotes millions and "th." thousands of DKK. The digitalisation index captures firms' use of digital technologies with values lying between $0 - 18$, and cloud importance is an indicator equal to one if cloud computing is reported as important for firm operations.

the advantage of linking survey responses to comprehensive administrative data, which allows us to benchmark survey participants against the universe of Danish firms.

2.2 Definitions and Descriptive Evidence on AI Adoption

Following the OECD, we define artificial intelligence (AI) as machine-based systems that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments (OECD 2019). This definition emphasises systems that rely on data-driven learning or algorithmic decision-making rather than rule-based automation.

In line with this definition, we define AI adoption as the use of at least one AI-related technology in production.[2] Our baseline empirical focus is on a set of *Core AI* technologies that most directly embody this definition and whose adoption reflects discrete, deliberate firm-level investment decisions. Concentrating on this margin generates meaningful cross-firm variation, while limiting measurement noise from broader digital investments that may predate the period of analysis.

To capture heterogeneity in technological complexity, functionality, and adoption costs, we distinguish between three categories of technologies covered by the survey:

- **Core AI**: Machine Learning (ML), Natural Language Processing (NLP), Machine Vision (MV), and Voice Recognition Software (VRS);
- **Generative AI**: large language models and related generative applications, explicitly added in the 2024 wave;
- **Other advanced technologies**: Robotics, Automated Guided Vehicles (AGV), Automated Storage Systems (ASS), Radio Frequency Identification (RFID), Touch Screens/Kiosks (TS), Augmented Reality (AR), and Automated Decision-Making Systems (ADMS).

Figure 1: Adoption of Core AI Technologies

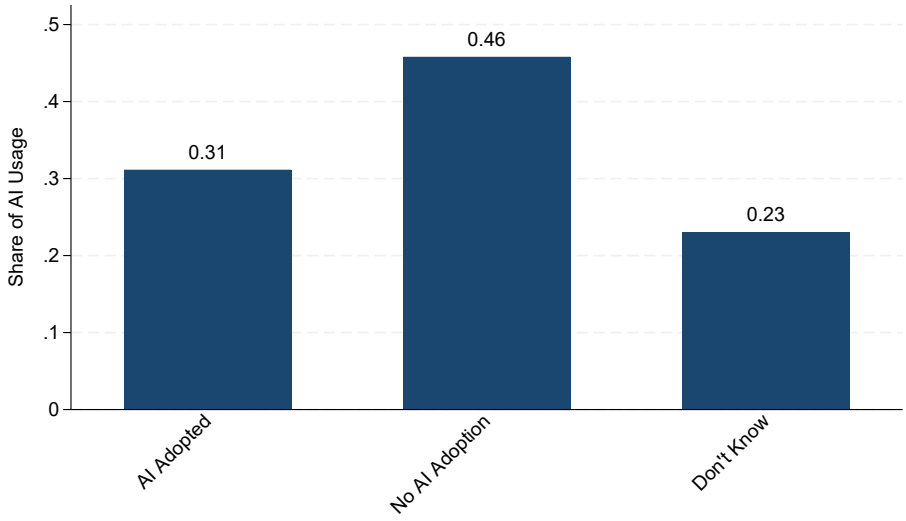

Note: The figure shows the weighted adoption rate of Core AI technologies among respondents to the survey. "AI adopted" denotes the share of firms that have implemented at least one Core AI technology (Machine Learning, Natural Language Processing, Machine Vision, and Voice Recognition Software). Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

The third group does not consist exclusively of AI technologies; however, many of its components are *AI-enabled*, such as robotics integrated with machine vision or automated

2. See Section A.2 in the Appendix for a detailed definition of the different technologies included in the survey.

decision-making systems powered by machine learning. We therefore interpret this group as a complementary set of advanced digital technologies that may interact with, or depend on, AI. We note that the boundaries between these categories are not fixed. Generative AI systems increasingly incorporate capabilities that overlap with core AI (e.g., NLP), and some technologies classified as "other" – notably automated decision-making systems – are now commonly built on machine learning. Our classification reflects the survey instrument and the technologies available at the time of data collection, not a claim that these distinctions will remain stable. Throughout the analysis, references to AI adoption correspond to the use of *Core AI* technologies unless stated otherwise. Adoption of the broader set of advanced technologies, as well as generative AI, is discussed for comparison.

Figure 2: Levels of AI Technology Adoption

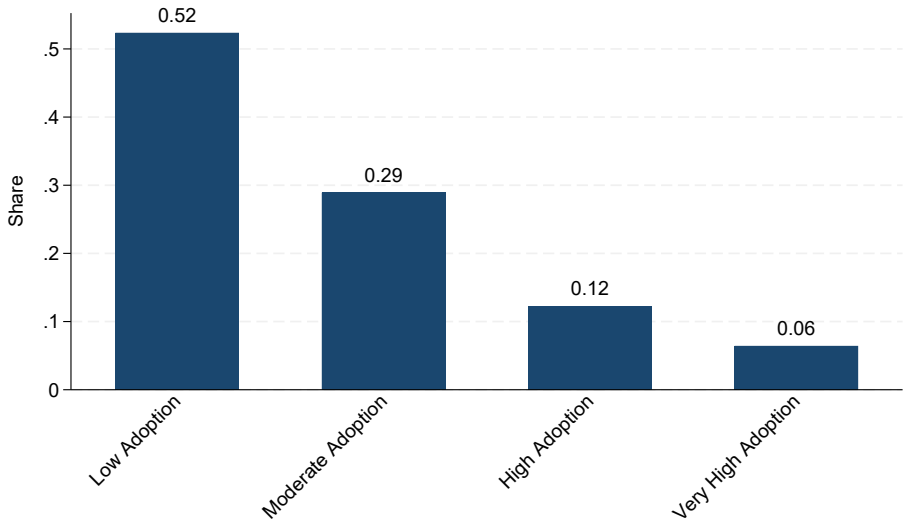

Note: The figure illustrates the distribution of AI adoption intensity among firms adopting at least one Core AI technology (Machine Learning, Natural Language Processing, Machine Vision, and Voice Recognition Software). Adoption levels correspond to the number of Core AI technologies adopted: Low (one), Moderate (two), High (three), and Very High (all four). Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Figure 1 presents the weighted adoption rate of Core AI technologies. Approximately 31% of firms report having adopted at least one Core AI technology. The figure also shows that a non-negligible (23%) share of respondents report being unsure about whether Core AI has been adopted, highlighting the informational frictions surrounding AI use within firms. In all analyses, firms reporting "don't know" are classified as non-adopters. This is a conservative choice: if some uncertain firms are in fact adopters, our adoption rates are lower bounds and estimated associations may be attenuated. We discuss robustness to alternative classifications in Section 6. Figure E.1.1 in the Appendix reports weighted adoption rates across the full set of technologies covered by the survey. Machine learning and natural language processing are the most widely adopted *Core AI* technologies (18-19%). By contrast, generative AI, which is included only in the 2024 survey wave, exhibits a weighted adoption rate of approximately 50% among all surveyed firms in that

wave. This notably higher rate likely reflects the lower barriers to generative AI adoption: whereas core AI typically requires deliberate firm-level investment in infrastructure and technical expertise, generative AI use may capture a broader range of behaviours, including individual employees using publicly available tools such as ChatGPT.

Figure 3: Adoption of AI Technologies VS Other Advanced Technologies by Industry

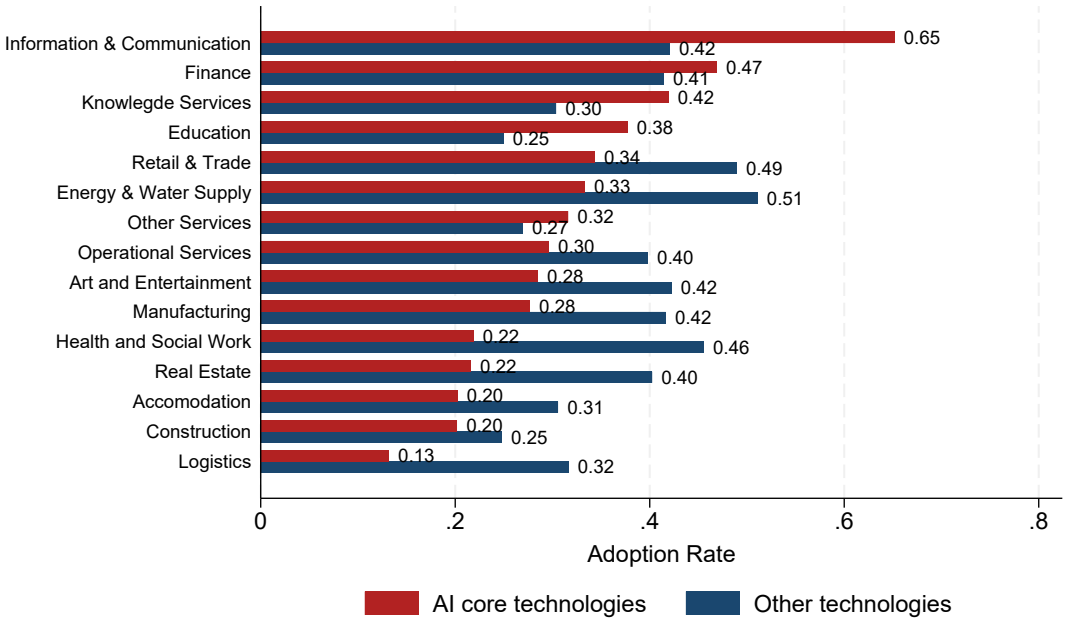

Note: This figure illustrates the weighted adoption rate of Core AI VS other technologies by industry. Core AI technologies are: Machine Learning, Natural Language Processing, Machine Vision, and Voice Recognition Software. Other technologies are: Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Beyond the extensive margin of *Core AI* adoption, Figure 2 examines the intensity of AI use among adopting firms. Restricting the sample to firms that report using at least one *Core AI* technology, the figure shows their distribution by the number of technologies deployed. Approximately half of adopters report using only one AI technology, while the remainder use two or more *Core AI* technologies. Such evidence points to AI adoption possibly occurring as part of a broader pattern of technological upgrading rather than in isolation. Appendix Figure B.1 reports a Jaccard co-adoption matrix across all technologies covered by the survey. The matrix shows that Core AI technologies are frequently adopted jointly, with particularly strong co-adoption between Natural Language Processing and Machine Learning (Jaccard index of 0.41), as well as between Machine Vision and Machine Learning (0.29). At the same time, the matrix reveals non-trivial complementarities between Core AI and selected advanced technologies, such as robotics and machine learning (0.28), indicating that AI adoption often coincides with broader investments in advanced production technologies.

Adoption is prevalent across all firm-size categories, but with distinct patterns for Core AI versus other advanced technologies. Core AI adoption is highest among the smallest firms, where around 40% report use, whereas firms with more than 100 employees have adoption at roughly 20%. By contrast, adoption of other advanced technologies is broadly stable across the size distribution, ranging between 31% and 41%. This divergence suggests that while robotics and related technologies diffuse more evenly across firm sizes, the relatively high prevalence of Core AI among smaller firms may partly reflect the presence of start-ups and smaller firms in knowledge-intensive business services (KIBS), where AI-based tools are closely aligned with digital service production. At the same time, this pattern may also suggest that some AI technologies involve lower fixed costs and greater scalability than capital-intensive technologies.

Figure 4: Adoption of AI Technologies VS Other Advanced Technologies by Firm Size

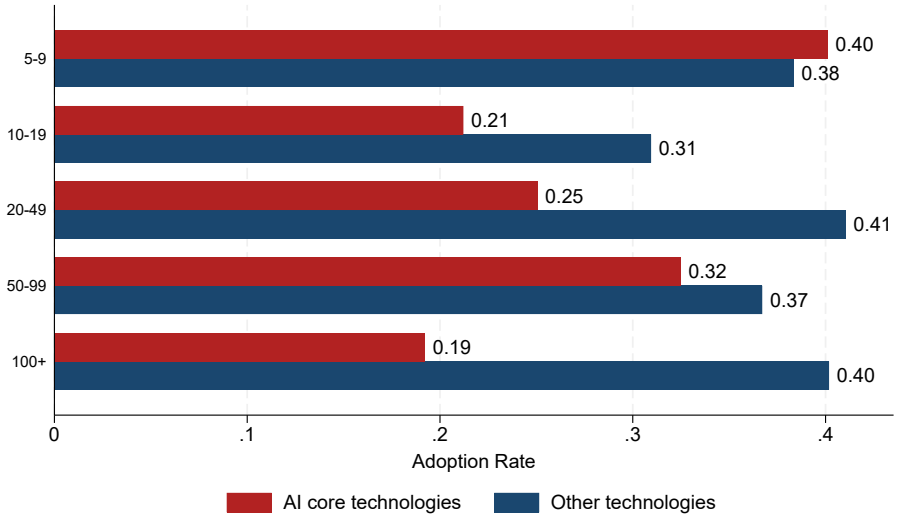

Note: This figure illustrates the weighted adoption rate of Core AI VS other technologies by firm size. Core AI technologies are: Machine Learning, Natural Language Processing, Machine Vision, and Voice Recognition Software. Other technologies are: Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Regarding industry heterogeneity, Figure 3 highlights that sectors rely on different technological bundles, reflecting differences in production processes and task composition. Core AI adoption is particularly concentrated in knowledge-intensive service sectors, including information and communication, finance, and professional services, while manufacturing and retail trade exhibit higher adoption of other advanced technologies.

Further descriptive evidence is provided by additional figures in the Appendix. For firms adopting at least one Core AI technology, Appendix Figure E.1.3 distinguishes between internally developed and externally acquired AI solutions by firm size. Among smaller firms (5–9 and 10–29 employees) and medium-sized firms with up to 100 employees, internal development accounts for roughly 30–37% of AI adoption, whereas among firms

with more than 100 employees the share exceeds 50%. This pattern highlights systematic differences in in-house capabilities and sourcing strategies across firm-size classes. The survey does not distinguish between building a system from scratch and developing an application on top of an existing platform, so internal development should be interpreted broadly. Appendix Figure E.1.4 presents the distribution of the digitalisation index, revealing substantial heterogeneity in the digital foundations supporting AI use.

Appendix Figure E.1.5 shows the share of employees with AI-related skills by firm size. Across all size classes, Core AI adopters exhibit a higher share of such employees than non-adopters, and the share increases monotonically with firm size. Among small and medium-sized adopting firms, about 41% of employees possess AI-related skills, rising to roughly 61% in firms with more than 100 employees. The direction of this association is a priori ambiguous: firms with AI-skilled workers may be better positioned to adopt, but firms planning to adopt may also recruit for AI skills in anticipation. Finally, Appendix Figure E.1.6 shows that concerns about AI adoption are concentrated around legal issues (45%) and insufficient infrastructure (40%), with acquisition costs (29%) representing a secondary barrier. Reputation risks (17%) and other concerns (13%) are mentioned less often. This distribution suggests that the diffusion of AI is constrained primarily by organizational and regulatory bottlenecks rather than by generalized resistance to the technology.

Table 2 provides an overview of a broad set of firm-level characteristics, comparing AI adopters, firms adopting other advanced technologies only, and non-adopters. Several systematic differences emerge. AI-adopting firms are substantially larger in terms of employment, revenues, and value added, and are more likely to engage in international trade. They also pay higher wages—both annually and hourly—and employ a more highly educated workforce, including a larger share of university graduates. Consistent with this, AI adopters have a more white-collar–intensive workforce composition than either other-technology adopters or non-adopters.

At the same time, AI adopters are not systematically different than other firms regarding firm age but exhibit stronger international orientation and higher capital intensity. These patterns suggest that AI adoption is associated with greater resource capacity, higher skill intensity, and stronger organisational capabilities. In contrast, firms adopting other advanced technologies without AI lie between AI adopters and non-adopters along most dimensions: larger and more international than non-adopters, but smaller and less skill-intensive than AI adopters.

Table 2: Summary Statistics by Adoption Type

	AI adopters (n=1,396)		**Other tech. adop. (n=1,733)**		**AI non-adop. (n=1,619)**		**Tech. non-adop. (n=2,374)**	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
No. of employees	63.05	225.59	60.14	206.22	25.32	54.10	21.59	41.52
Revenue (M)	323.29	5336.66	222.53	1277.07	66.64	305.88	51.04	198.12
Value added (M)	69.22	439.00	58.01	328.61	19.79	51.53	16.38	43.08
Yearly wage (th.)	414.44	185.59	389.56	186.11	351.64	159.22	346.81	153.25
Hourly wage	266.80	92.98	253.98	89.60	234.45	72.40	232.28	70.31
Firm age	18.11	16.65	20.08	17.45	18.81	15.55	18.12	15.06
Exports (M)	152.96	3873.18	79.91	919.64	16.02	112.94	10.51	81.89
Imports (M)	244.99	8614.87	148.02	6486.19	15.16	109.07	10.30	72.96
Physical capital stock (M)	19.92	289.78	16.27	257.24	4.49	44.00	3.78	43.85
Intangible capital stock (M)	0.05	0.97	0.04	0.89	0.02	0.24	0.03	0.26
Foreign ownership	0.16	0.37	0.15	0.36	0.09	0.28	0.08	0.27
Digitalisation index	13.71	4.25	12.85	4.32	10.81	4.82	10.42	4.95
Cloud importance	0.88	0.33	0.81	0.39	0.70	0.46	0.67	0.47
Female workers	0.36	0.25	0.37	0.27	0.36	0.28	0.35	0.28
Immigrant workers	0.16	0.21	0.14	0.20	0.13	0.20	0.13	0.21
Avg. age (workers)	40.57	7.22	41.77	7.62	42.05	8.33	41.84	8.42
Younger than 36	0.42	0.24	0.38	0.24	0.38	0.25	0.38	0.25
Between 36 and 54	0.39	0.18	0.39	0.17	0.38	0.18	0.38	0.19
Older than 55	0.19	0.17	0.23	0.18	0.25	0.20	0.24	0.20
No secondary education	0.14	0.17	0.18	0.18	0.21	0.19	0.22	0.19
Bachelor's or higher	0.47	0.31	0.36	0.28	0.28	0.27	0.27	0.27
Master's or higher	0.21	0.23	0.13	0.18	0.09	0.16	0.08	0.16
STEM educated	0.24	0.25	0.24	0.25	0.24	0.26	0.24	0.26
University STEM educated	0.14	0.20	0.10	0.17	0.06	0.14	0.06	0.13
Union membership	0.59	0.22	0.61	0.22	0.58	0.23	0.57	0.23
Managers	0.07	0.10	0.07	0.10	0.06	0.09	0.06	0.09
Professionals	0.46	0.35	0.36	0.33	0.27	0.32	0.26	0.32
Clerical/Administrative	0.13	0.19	0.12	0.18	0.11	0.17	0.11	0.16
Sales/Services	0.10	0.21	0.11	0.21	0.14	0.24	0.15	0.25
Skilled trades	0.11	0.24	0.13	0.26	0.17	0.29	0.18	0.30
Production workers	0.13	0.23	0.20	0.28	0.23	0.31	0.23	0.31

Notes: The table reports means and standard deviations of firm-level variables of respondents to the survey by adoption status. AI adopters use at least one core AI technology (machine learning, natural language processing, voice recognition, machine vision). Other tech. adopters use advanced non-AI technologies but no core AI. Other technologies are: Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems. AI non-adopters do not adopt core AI (and may adopt other advanced technologies). Tech. non-adopters adopt neither. All statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Monetary values are nominal; "M" denotes millions and "th." thousands of DKK. Occupational groups are shares of total employment.

3 Predictors of AI Adoption

The descriptive evidence in Section 2.2 documents substantial heterogeneity in AI adoption across firms, industries, and regions. These patterns suggest that adoption reflects systematic differences in firms' technological capabilities, organisational structure, and workforce composition. In this section, we formalise these relationships using regression analysis that combines the AI adoption survey with rich firm-level and linked employer–employee administrative data.

First, we study the extensive margin of AI adoption, examining which predetermined firm characteristics are associated with the adoption of any AI technology (i.e., whether a firm adopts at all). Second, conditional on adoption, we analyse the intensive margin of advanced technological use, identifying which firm characteristics predict the breadth of technological use (i.e., the number of distinct technologies deployed).

Finally, we examine how AI is used within firms by documenting heterogeneity of technology use across business functions. This allows us to investigate whether a distinction emerges between general-purpose technologies that diffuse broadly across organisational activities and more specialised technologies whose use is concentrated in particular domains.

3.1 Empirical Approach

The objective of the following analysis is to assess the relative importance of firm-level capacities and workforce characteristics as predictors of both the decision to adopt AI and the breadth of technological deployment. We identify firm-level AI adoption using survey responses collected in 2023 and 2024, while firm characteristics, productivity measures, and workforce composition from administrative registers are measured in 2022. This timing mitigates concerns of simultaneity and reverse causality for the register-based variables. Two covariates, however, are contemporaneous: firms' digitalisation index and cloud storage importance are obtained from the survey itself and therefore reflect digital readiness at the time of the adoption decision rather than a pre-determined state. Because these are among the strongest predictors in our specifications, the associations involving these variables should be interpreted with particular caution.

3.1.1 Extensive Margin

To quantify the firm-level correlates of AI adoption, we estimate the following regression model:

$$\Pr(AI_j = 1) = \Lambda\left(Z_j'\delta + X_j'\beta + \theta_s + \gamma_p\right) \tag{1}$$

where $\Lambda(\cdot)$ denotes the logistic function, AI_j is a binary indicator equal to one if firm j adopted at least one Core AI technology in the 2023 or 2024 survey waves.[3] The vector Z_j captures firm-level characteristics, X_j describes workforce composition, both measured in 2022. Sector and province fixed effects are denoted by θ_s and γ_p, respectively. Some specifications additionally control for unobserved firm heterogeneity by including ventiles of the firm-level total factor productivity distribution, estimated over the period 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015.[4] All regressions are weighted by the inverse probability of responding to the survey, and standard errors are robust to heteroskedasticity.[5] In addition to core AI adoption, we estimate parallel models for the adoption of other advanced technologies, as well as for generative AI (measured only in the 2024 wave), to benchmark AI adoption against broader patterns of technology adoption.

Drawing on the innovation literature and technology adoption literature, we include firm-level covariates Z_j and X_j capturing internal capabilities and human capital associated with technology adoption and innovation performance. Firm size, measured by the logarithm of total employment, proxies for scale and absorptive capacity: larger firms may be better able to bear fixed costs of adoption and implement organisational change (Cohen, Levinthal, et al. 1990; McElheran et al. 2024). Firm productivity, measured by logarithm of value added per worker, helps separate scale effects from efficiency (Bloom et al. 2012). Capital intensity and intangible assets per worker account for complementary investments that may enhance the returns to AI (Lee et al. 2022). Exposure to external knowledge and technological frontiers may further facilitate adoption. Firm age is included to capture potential organisational rigidities: younger firms may be more flexible and more likely to adopt emerging technologies (McElheran et al. 2024). A foreign ownership indicator reflects exposure to global technological frontiers and advanced management practices (Bloom and Van Reenen 2007), while exporter and importer dummies proxy for access to international markets that may facilitate technology diffusion (Keller 2004; Bloom et al. 2016).

The log of average hourly wages serves as a proxy for workforce quality which may complement AI technology, even conditioning on detailed worker characteristics. Worker representation may also influence the pace of technological adoption. Union presence can affect firms' incentives to introduce labour-saving technologies through wage bargaining and negotiations over organisational change, although it may also facilitate coordination

3. "Don't know" responses to the adoption question are kept in the sample and coded as non-adopters. This may attenuate estimated coefficients if some of these respondents are actual adopters, implying that results should be interpreted as lower-bound estimates.

4. The TFP sample of firms is smaller than the main estimation sample due to missing information on inputs required for TFP estimation or because the firm was not active before 2015. Table F.1 shows that the composition of the sample for this robustness check is not radically different from the main sample.

5. Appendix C reports details of the estimation of the survey response probability.

and information sharing during technological restructuring (Tauman and Weiss 1987; Kostøl and Svarstad 2023; Bryson and Forth 2023). To capture variation in bargaining power across firms, we include the share of workers who belong to a union and an indicator for whether the firm belongs to an employer association.

We incorporate direct measures of digital infrastructure which can be regarded as preconditions for effective AI deployment: the extent of data digitalisation and the use of cloud computing services.[6] This builds on a broader literature that highlights complementarities between digital infrastructure, organisational change, and the adoption of advanced technologies (Bresnahan et al. 2002; Goldfarb and Tucker 2019).[7]

Human capital and workforce composition may be associated with firms' ability to adopt and integrate new technologies. Educational composition is measured by the shares of workers with basic, secondary, and tertiary education, together with the STEM share, reflecting technical absorptive capacity. Higher education and technical skills are likely complementary to AI technologies, which often require adaptation, monitoring, and integration into existing tasks. We further include the shares of younger (below 36) and older employees (above 55) to capture differences in openness to new technologies and the risk of skill depreciation. The shares of white-collar employees and managers proxy for organisational hierarchy and knowledge intensity, which may be associated with the scope and coordination of AI use.[8] Drawing on evidence linking workforce structure and diversity to innovation and productivity (Parrotta et al. 2014b; Parrotta et al. 2014a; Iranzo et al. 2008; Navon 2009; Østergaard et al. 2011; Hong and Page 2004), we include a set of controls capturing the demographic composition of each firm's workforce: the shares of female and foreign-born workers. Differences in workforce backgrounds may broaden the range of perspectives and problem-solving abilities within the firm, although demographic diversity may also generate coordination and communication costs and potential trust frictions within teams (Hong and Page 2004; Alesina and La Ferrara 2005).

We verified that multicollinearity is not a concern in our specifications. All variance inflation factors (VIFs) remain well below the conventional threshold of 5, with the highest

6. Digitalisation is measured using the survey question on the extent to which information in six business functions (HR, finance, marketing, customer feedback, supply chain, and production) is stored in digital format. We construct an index ranging from 0 to 18 and standardise it for regression analysis. Figure E.1.4 shows the distribution of digitalisation intensity across firms. Cloud usage is an indicator equal to one if respondents rate cloud services as "pretty important" or "very important" for the firm's IT functions. Note that some cloud services increasingly bundle AI capabilities (e.g., AI-powered analytics or coding assistants), so cloud importance and AI adoption may partly reflect the same underlying technological choice.

7. Relevant discussions include absorptive capacity (Cohen, Levinthal, et al. 1990), management practices and productivity (Bloom and Van Reenen 2007; Bloom et al. 2012), and complementarities between ICT, organisational change, and intangible capital (Bresnahan et al. 2002; Corrado et al. 2009; Goldfarb and Tucker 2019).

8. See also Parrotta et al. 2014b; Autor et al. 2003; Acemoglu and Autor 2011; Deming 2017; Garicano and Rossi-Hansberg 2006 for additional discussion of how workforce composition and organisational structure affect innovation and technological adoption.

VIF equal to 3.36. Appendix D implements an alternative, data-driven empirical strategy that is not sensitive to an arbitrary ordering of covariates. Using a lasso-logit model, we select the most relevant predictors of AI adoption from a broader set of potential covariates based on their predictive power, and we assess the stability of the main correlates identified in the baseline specification.

3.1.2 Intensive Margin

While the extensive margin captures which firms adopt AI at all, it does not provide information on the breadth of technology use within the firm. Firms that adopt AI may differ substantially in the number of distinct technologies they deploy, reflecting heterogeneity in their ability to integrate these technologies across multiple activities. Analyzing the intensive margin therefore allows us to assess whether the predictors of AI adoption differ between marginal adopters and firms with more extensive technological use. While initial adoption may be predicted by broad firm characteristics such as size, sectoral exposure, or baseline digital infrastructure, the depth of adoption may be more strongly associated with complementary organisational capabilities and workforce skills.

We assume that the number of technologies adopted by firm j, Y_j, follows a zero-truncated Poisson distribution with the conditional mean parameter specified as follows:

$$\lambda_j = \exp\left(Z_j'\delta + X_j'\beta + \theta_s + \gamma_p\right) \tag{2}$$

where Z_j and X_j denote firm and workforce characteristics of firm j, with covariates measured in 2022, and θ_s and γ_p denote sector and province fixed effects.[9] The baseline set of covariates is the same as in the extensive-margin model. For comparison, we also estimate the same specification on samples of firms that have adopted other innovations or generative AI, to examine whether the predictors of the breadth of technology adoption differ across firms adopting AI, generative AI, or other digital innovations.

3.1.3 AI Adoption Across Business Functions

To understand how AI is deployed within firms, we analyse adoption across business functions. Specifically, we ask: conditional on adopting a given technology k, how does the likelihood of using AI differ across business functions f? To answer the question, we combine information from survey questions on the types of technologies adopted in the firm and the business functions where AI is used. We estimate the following logit model:

9. As in the extensive-margin model, some specifications additionally control for unobserved firm-specific heterogeneity by including ventiles of the total factor productivity (TFP) distribution, estimated over 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015. All regressions are weighted by the inverse probability of responding to the survey, estimated using the approach described in Appendix C. Standard errors are robust to heteroskedasticity.

$$\Pr(AI_{jf} = 1) = \Lambda\left(\alpha_f + \sum_{k=1}^{K} \gamma_k \, Tech_{jk} + \sum_{k=1}^{K}\sum_{f=1}^{F} \delta_{fk} \left(Tech_{jk} \times \mathbf{1}\{F_j = f\}\right) + X_j'\beta\right) \quad (3)$$

where $\Lambda(\cdot)$ denotes the logistic function. The dependent variable equals one if firm j reports using AI in business function f. The indicator $Tech_{jk}$ denotes whether firm j adopts technology k, while F_j denotes the business function. Function fixed effects α_f absorb baseline differences in AI use across business functions, while the coefficients of interests are γ_k and δ_{fk}, which give the marginal effect of adopting technology k on the probability of reporting AI use in business function f. The vector X_j includes firm-level controls: firm size, sector, and province fixed effects. The specification is estimated on the sample of firms that adopt at least one technology k and weighted by the inverse probability of responding to the survey, estimated with the approach described in Appendix C. Standard errors are robust to heteroskedasticity.

3.2 Who Adopts AI? Main Findings

In this section we discuss the main results on technology adoption patterns. We begin with the firm-level correlates of adoption, who adopts AI and how broadly, before turning to how adopters deploy AI across business functions.

3.2.1 Extensive Margin

Table F.2 reports average marginal effects from logit models estimating the probability that a firm adopts advanced technologies. We distinguish between adoption of Core AI technologies, other advanced technologies, and Generative AI. Firm-level characteristics alone explain a substantial share of variation in Core AI adoption, with a baseline Pseudo R^2 of 0.11. Adding workforce composition in terms of education, age and occupation raises it to 0.15, pointing to an important role for human capital and workforce structure. Including sector and province fixed effects increases the Pseudo R^2 further to 0.167, and adding further workforce diversity measures brings it to 0.171. Finally, including firm TFP ventiles raises it slightly to 0.174. Overall, firm characteristics and workforce composition account for most of the explained heterogeneity in adoption, while proxies for productivity adds relatively little beyond them.

Figure 5 illustrates estimates from the specification including firm characteristics, workforce composition, and sector and province fixed effects, corresponding to columns (3), (8), and (14) in Table F.2. Across all technology categories, firm size emerges as a key and robust predictor of adoption. For Core AI, a 10 percent increase in employment is associated with an increase of roughly 0.5 percentage points in the probability of adoption.

Figure 5: Predictors of Technology Adoption - Extensive Margin

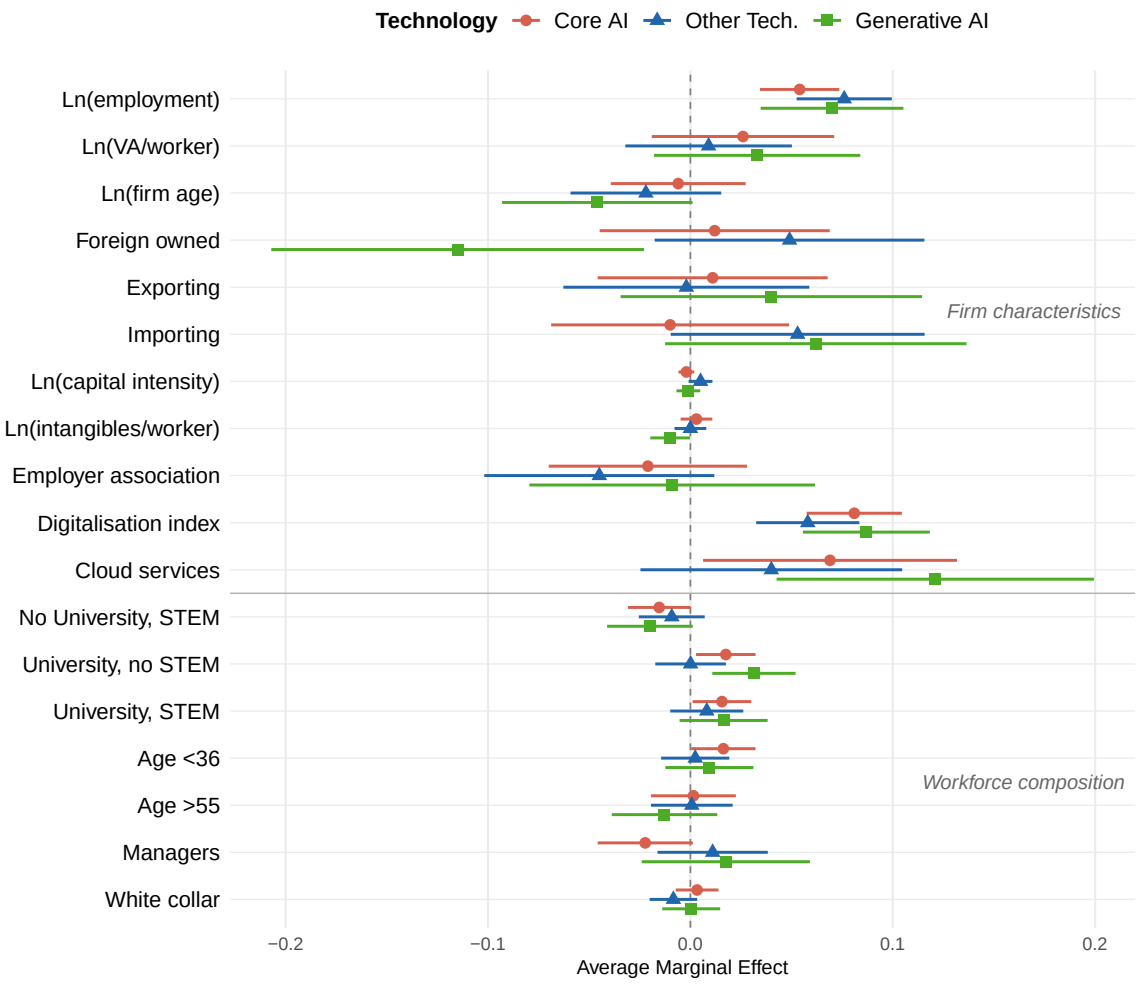

Note: The figure displays average marginal effects (AMEs) from logit models of technology adoption at the extensive margin, estimated separately for three technology categories: Core AI (Machine Learning, Natural Language Processing, Machine Vision, or Voice Recognition Software), Other Technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems), and Generative AI. Estimates correspond to columns (3), (8), and (14) in Table F.2, that is, specifications including firm characteristics, workforce composition, and sector and province fixed effects. Whiskers denote 95% confidence intervals based on robust standard errors. Workforce composition variables (education, age, occupational, and gender shares) enter the regression as shares bounded between zero and one; for visual comparability with other variables, their AMEs and confidence intervals are rescaled to reflect a 10 percentage point change in the relevant share. Exporting, Importing, Foreign owned and Cloud services are binary indicator variables. The digitalisation index is standardised. Regressions are weighted by the inverse probability of responding to the survey, estimated as reported in Appendix C.

These magnitudes are economically meaningful but moderate, indicating that scale effects operate gradually rather than reflecting large discrete adoption thresholds. The relationship is similar for other advanced technologies and stronger for Generative AI, suggesting that even for software-based AI applications, organisational scale remains an important facilitating factor. This pattern is consistent with evidence on scale economies and complementary investments in technology adoption (Bresnahan et al. 2002; Brynjolfsson and Hitt 2000).

Digital readiness is a central and quantitatively important predictor. A one-standard-deviation increase in the digitalisation index is associated with an 8–10 percentage point

higher probability of adopting core AI technologies. A similar pattern emerges across all three technology groups, with coefficients that remain among the largest and most stable across specifications. Cloud use reinforces this interpretation: firms relying on cloud-based infrastructure are substantially more likely to adopt Core AI and Generative AI technologies, while it has no predictive power for adoption of other advanced technologies. This contrast points to heterogeneity across technology types, with AI adoption appearing to depend more strongly on digital infrastructure than the adoption of other advanced technologies.

International exposure is associated with adoption in more selective ways. Importing firms are more likely to adopt other advanced technologies and Generative AI in several specifications, although these effects are not uniformly robust across all models. Exporting status, by contrast, is generally not significantly related to adoption once other firm characteristics are controlled for. This pattern is consistent with the view that global sourcing can facilitate access to new technologies and embedded knowledge (Keller 2004; Bloom et al. 2016), but does not support a strong or systematic role for international exposure as a primary driver of adoption.

In contrast, standard firm performance measures have limited predictive power once sector and regional heterogeneity is taken into account. Value added per worker is generally small and statistically insignificant across technologies, suggesting that adoption is not mechanically driven by productivity differences. Consistent with this, adding ventiles of firm TFP distribution in the last column of each panel leads only to a marginal increase in the Pseudo R^2, while leaving the main patterns for the other predictors unchanged. This suggests that latent productivity has limited explanatory power for adoption beyond the rich set of observed firm characteristics.

Capital intensity and intangible intensity are also generally small and insignificant across technologies, implying that adoption is not simply a reflection of capital deepening. Instead, these results point to differences in organisational readiness and complementary investments. Firms with similar observable performance and capital profiles may therefore face very different adoption prospects, consistent with views of heterogeneous technology diffusion (Brynjolfsson and Hitt 2000). Institutional measures also play only a limited role. Employer association membership is not robustly associated with adoption once richer controls are included, and the share of union members in the workforce is likewise not a systematic predictor across technologies.

Workforce composition is an important predictor of AI adoption. More precisely, the education-share variables indicate that firms with a higher share of workers in the *University, no STEM* category are significantly more likely to adopt Generative AI, and Core AI technologies. The share of workers in the *University, STEM* category is also positively associated with Core AI adoption. However, this relationship becomes non-significant once controls for female and immigrant shares are included, indicating

that differences in STEM composition across firms are correlated with demographic composition. For example, firms with a higher share of university STEM workers tend also to have systematically lower share of female employees. By contrast, firms with a higher share of workers in the *No University, STEM* category are less likely to adopt Core AI and Generative AI. This suggests that, beyond field of specialization, the level of formal education matters for adoption, with tertiary-educated STEM workers appearing more relevant complements to these technologies than workers with STEM backgrounds but no university degree. Other workforce characteristics play a more nuanced role. In particular, occupational composition variables do not exhibit strong predictive power once education measures are included. Also age and gender composition show limited and unstable associations across specifications. For Generative AI, firms with higher shares of immigrant workers are significantly less likely to adopt, even within sectors and regions.

To address potential concerns that the estimates in Table F.2 are influenced by collinearity or by the sequence in which covariates are introduced, Appendix D implements a lasso-logit model that selects the most relevant predictors of AI adoption from a broader set of candidate covariates based on their predictive power. This alternative, data-driven strategy is not affected by an arbitrary ordering of covariates and shows that the main correlates identified in the baseline specification remain stable. Tables D.1 and D.2 confirm that firm size and digitalisation are key predictors of adoption across all advanced technologies, while cloud-service use is particularly associated with the adoption of Core AI and Generative AI. In addition, the baseline share of the workforce in the *University, STEM* category also predicts adoption of Core AI technologies.

We further examine whether the correlates of AI adoption vary systematically across industries by estimating sector-specific adoption models for manufacturing, trade and transport, knowledge-intensive business services (KIBS), and other services. Appendix Table F.3 shows that digitalisation and the share of university STEM-educated workers significantly predict adoption in manufacturing and KIBS, whereas they are not significant in trade and transport. Similarly, the use of cloud services significantly predicts adoption of Core AI primarily in the KIBS sector.

Overall, the results paint a clear picture of which firms adopt advanced technologies. Conditional on sector and region, adoption is concentrated among larger firms with higher digital maturity, access to cloud-based infrastructure, and workforces with larger shares of tertiary-educated workers, particularly those with STEM degrees for adoption of Core AI and non-STEM degrees for the adoption of Generative AI. By contrast, firms that are smaller, less digitally prepared, or lack complementary human capital are substantially less likely to adopt, even when operating in similar environments. Importantly, the strength and nature of these predictors differ across technology types: while firm size and digital infrastructure are common drivers, workforce composition plays a particularly prominent role for Core AI and Generative AI. To sum up, the extensive-margin findings indicate

that AI adoption is technology-specific, and associated with distinct complementary organisational and workforce capabilities rather than diffusing uniformly across firms.

3.2.2 Breadth of Adoption: Intensive Margin Results

Table F.4 reports average marginal effects from truncated Poisson models for the intensive margin of technology use. In all columns, the dependent variable is the number of distinct technologies adopted from the full list of surveyed technologies. The intensive-margin analysis is conditional on adoption of at least one technology within each category: columns (1)–(5) restrict the sample to firms that have adopted at least one Core AI technology; columns (6)–(10) to firms that have adopted at least one Other Technology; and columns (11)–(15) to firms that have adopted Generative AI. The estimates, therefore, capture how firm characteristics relate to the breadth of overall technology portfolios among adopters, rather than the decision to adopt a particular technology category. The table reports average marginal effects, therefore coefficients are interpreted in technology-count units.

A comparison of the extensive and intensive margin results reveals that crossing the adoption threshold and scaling AI within the firm are associated with partially distinct sets of firm characteristics. Firm size and the digitalisation index are robust predictors at both margins, confirming that scale and digital maturity facilitate adoption throughout the diffusion process. Firm size is the most consistent predictor of more intensive technology use across all adopter samples. Among Core AI adopters, a 10 percent increase in employment is associated with approximately 0.03–0.04 additional technologies adopted. The corresponding magnitude is similar for adopters of other advanced technologies and somewhat larger for Generative AI adopters, where the increase is around 0.04–0.05 technologies.

Digital maturity remains an important predictor. A higher digitalisation index is consistently associated with the adoption of a broader set of technologies across all adopter groups, with particularly large effects for other advanced technologies (approximately 0.45–0.50 additional technologies for a one-standard-deviation increase). This suggests that firms that initially adopt other advanced technologies, such as robotics or automated storage systems, have invested in function-specific physical systems, and that broadening their technology portfolio depends more on digital infrastructure and intangible capital than it does for firms that have already crossed the threshold through AI adoption. For Core AI and Generative AI adopters, the effects are smaller but remain positive and statistically significant. Taken together with the extensive-margin results, these findings indicate that digital capabilities both facilitate adoption and support subsequent scaling.

Beyond these common drivers, however, the two margins diverge. Cloud services strongly predict initial adoption, particularly for Core AI and Generative AI, but show

no positive association with technology breadth among adopters, suggesting that cloud infrastructure lowers barriers to entry without necessarily supporting deeper deployment. Conversely, foreign ownership is weakly and inconsistently associated with adoption at the extensive margin but emerges as a strong and significant predictor of technology breadth among Core AI and other technology adopters. This pattern is consistent with multinationals providing the managerial practices and cross-border knowledge transfers that support firm-wide deployment.

A similar divergence emerges with respect to workforce composition. At the extensive margin, the type of human capital matters: educational composition, particularly the distinction between STEM and non-STEM university-educated workers, predicts whether a firm adopts Core AI or Generative AI. At the intensive margin, conditional on having adopted Core AI, what predicts broader deployment is not workforce composition but its overall quality, as captured by average hourly wages. Among Core AI adopters, firms with higher average wages tend to deploy a broader set of technologies: a 10 percent increase in average hourly wages is associated with approximately 0.15 additional technologies adopted. This distinction suggests that crossing the adoption threshold requires specific technical expertise, particularly STEM skills for Core AI, whereas, once a firm has crossed that threshold, the deployment of additional technologies across the organisation is associated with higher-quality workers, conditional on the level and field of education.

3.2.3 Heterogeneity by Business Functions

Adoption alone provides limited information about how AI reshapes firm activities. Even conditional on adoption, AI use is unlikely to be uniform within firms. Differences in task content, information intensity, and organisational scope across business functions can shape where AI technologies are ultimately used (Autor et al. 2003; Acemoglu and Restrepo 2020). This section, therefore, focuses on heterogeneity in AI deployment across business functions within firms, shedding light on how AI is integrated into different parts of the production and organisational process.

Using data on AI use across business functions, we estimate the logit model in equation (3) relating the adoption of specific AI technologies to the probability that AI is used in any of the following business functions: customer relationship management, supply chain management, finance, human resources, marketing, and production. Average marginal effects are reported in Appendix Table F.5 and summarised visually in Figure 6. All reported coefficients are interpreted as percentage-point differences in the probability that AI is used in a given business function when a firm adopts a specific technology, relative to otherwise similar firms that do not adopt that technology.[10]

10. The sample includes firms adopting at least one technology. The control variables include sector, province and size bin fixed effects. To maximise sample coverage and comparability across technologies, these specifications exclude Generative AI which is included only in the 2024 wave of the survey. Appendix

Figure 6: Deployment of AI Technologies by Business Function

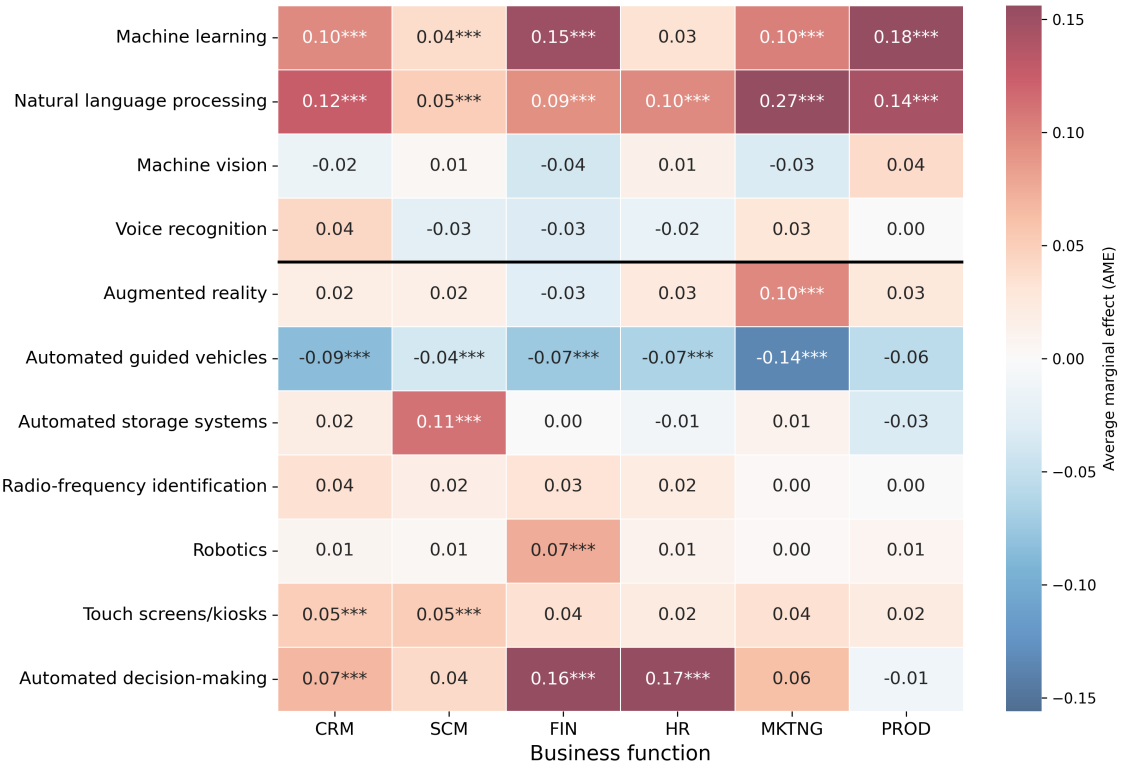

Note: The figure visualises average marginal effects from the logit model in equation 3, and computed from the estimated coefficients γ_k and δ_{fk}. The model relates the adoption of specific AI technologies in the firm to the probability that AI is used in a given business function. The 4 technologies above the black solid line are those classified as "Core AI" in section 2.2. The dependent variable is an indicator for AI use in the business function listed in the column, regressed on indicators for adoption of the technology listed in the row. Acronyms: CRM = Customer Relationship Management; SCM = Supply Chain Management; FIN = Finance; HR = Human Resources; MKTNG = Marketing; PROD = Production. The sample comprises adopters of at least one technology (N = 2,369 firms; 12,828 observations). All specifications control for firm size and include sector and province fixed effects. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Robust standard errors are reported in Table F.5. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The results reveal substantial heterogeneity in how different AI technologies are deployed across business functions and hint to a clear distinction between general-purpose and function-specific technologies. Among Core AI technologies, machine learning and natural language processing exhibit the broadest and most economically meaningful functional footprint. Holding constant adoption of other surveyed technologies and firm characteristics, firms that adopt machine learning are, on average, 15 percentage points more likely to report AI use in finance and around 18 percentage points more likely to report AI use in production than otherwise similar firms that do not adopt machine learning. Adoption of machine learning is also associated with increases of roughly 10 percentage points in the probability of AI use in marketing and customer relationship management. These magnitudes indicate that machine learning, once adopted, is integrated into both operational and decision-oriented functions within the firm.

Table F.6 and Figure F.1 report average marginal effects from analogous logit regressions estimated separately for each business area, using the 2024 subsample to assess how Generative AI adoption is associated with AI use across business functions.

Natural language processing displays an even more pervasive pattern of functional deployment. Firms adopting NLP are approximately 27 percentage points more likely to use AI in marketing, 15 percentage points more likely in production, and around 12 percentage points more likely in customer relationship management. In addition, NLP adoption is associated with increases of roughly 9–10 percentage points in the probability of AI use in finance and human resources. This broad functional reach suggests that language-based AI technologies diffuse widely across information-intensive, coordination-heavy, and decision-support functions.

By contrast, other Core AI technologies exhibit a much narrower and less consistent deployment pattern. For machine vision, estimated effects are small and statistically insignificant across all business functions, indicating that its adoption does not concentrate in specific business areas. Voice recognition shows similarly limited and imprecisely estimated associations. These patterns highlight that even within the Core AI category, technologies differ substantially in their scope for cross-functional deployment, reflecting differences in task specificity and integration requirements.

Other advanced technologies exhibit a markedly more function-specific deployment pattern, underscoring their close alignment with narrowly defined operational tasks rather than broader organisational processes. These patterns are consistent with prior evidence showing that automation technologies tend to complement specific workflows rather than enabling firm-wide organisational change (Graetz and Michaels 2018; Acemoglu and Restrepo 2020). Specifically, Automated storage systems are strongly concentrated in supply chain management: firms adopting these systems are 11 percentage points more likely to use AI in SCM, with no meaningful associations in other functions. Robotics adoption is primarily associated with finance-related activities, increasing the probability of AI use in finance by approximately 7 percentage points, consistent with its role in back-office and process-oriented applications. Touch screens and kiosks are associated with AI use in customer-facing and logistics-related functions, increasing the probability of AI use in customer relationship management and supply chain management by around 5 percentage points. Augmented reality is more prominent in marketing, suggesting that firms may use this technology for interactive product visualization or to enhance customer engagement. The negative coefficients for automated guided vehicles indicate that AI use is strongly concentrated in production-related activities, with systematically lower likelihoods of AI use in other business functions relative to adoption of more general-purpose AI technologies.[11] Finally, automated decision-making systems exhibit strong associations with finance and human resources: adoption of these systems is associated with increases of 16–17 percentage points in the probability of AI use in these functions, reflecting

11. Furthermore, relative to other technologies, firms may not perceive automated guided vehicles as AI, so their adoption is not consistently reported as AI use across business functions.

their close alignment with structured decision processes such as screening, allocation, and internal control.

Overall, the business-function results highlight that heterogeneity in AI use arises not only from differences in whether firms adopt AI, but also from how adopting firms deploy these technologies within the organisation. While some Core AI technologies, particularly machine learning and natural language processing, diffuse across multiple business functions, other AI and automation technologies remain tightly linked to specific operational domains. This distinction underscores the importance of accounting for the task orientation and organisational scope of technologies when assessing the implications of AI adoption for firm organisation and labour demand.

4 Individual-Level Evidence on AI Adoption and Workforce Readiness

The previous section showed that AI adoption among Danish firms is strongly associated with firm-level capabilities. Larger firms, firms with more advanced digital infrastructure, such as data digitalisation and cloud services, and those with more skill-intensive workforces are substantially more likely to adopt AI.

Firm-level adoption is measured through managerial responses and may therefore be subject to information frictions or reporting biases. AI use can emerge in decentralised or task-specific ways—for example, through experimentation within teams or through AI-enabled functionalities embedded in software—which may not always be fully visible to managers. At the same time, managers are often directly involved in technology adoption decisions and are therefore well positioned to observe and report firm-wide AI deployment, particularly when adoption reflects strategic or organisational changes.

To assess the reliability of managerial reporting and to provide a complementary individual-level perspective, we turn to worker-level evidence. We use a representative survey of approximately 4,000 individuals in Denmark, collected in two waves in 2023 and 2024. Respondents report whether AI is used at their current or most recent workplace, therefore capturing firm-level deployment rather than individual use. This allows us to benchmark worker-reported adoption against firm-level patterns and to examine whether worker-level reporting patterns across occupations and firm characteristics mirror the correlates identified in the firm-level analysis.

4.1 Self-Reported AI Use in the Workplace

Figure 7 shows that approximately one third of respondents report that AI is used in their current or most recent workplace, closely aligning with the firm-level adoption rate documented in Section 3 and Figure 1. A subset of respondents indicates that they

do not know whether AI is used; these responses are concentrated among non-employed individuals, who plausibly have weaker information about workplace technologies.

Figure 7: Worker-Reported Adoption of AI in the Workplace

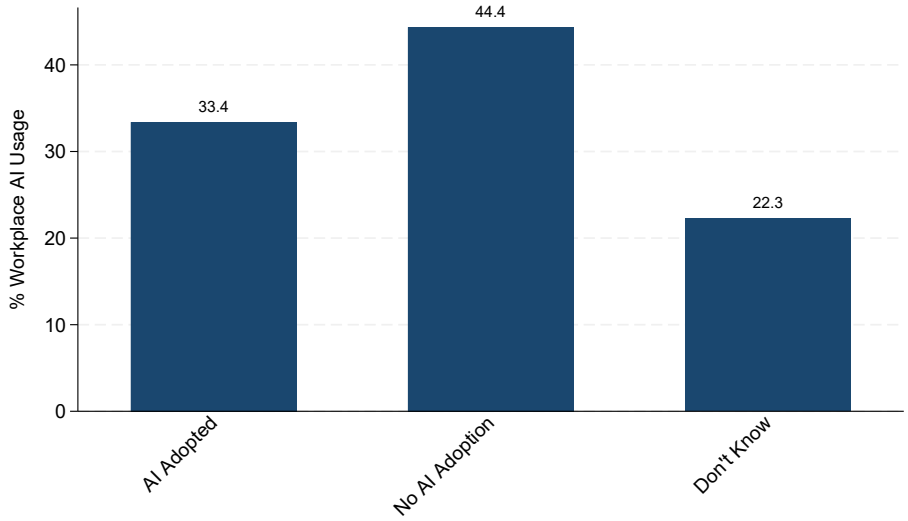

Note: The figure reports the share of individuals who report that the technology is used at their workplace. AI technologies are: machine learning, natural language processing, voice recognition, and machine vision. Statistics are weighted using survey weights.

Figure E.2.7 shows that the relative prevalence of technologies closely mirrors the firm-level evidence presented in Figure E.1.1. Established digital and automation technologies—such as robotics, tracking systems, and automated support tools—remain more widespread than AI applications. Within AI technologies, machine learning and natural language processing are reported more frequently than computer vision and related applications. The similarity in both ordering and relative prevalence across worker- and firm-level surveys suggests that worker reports capture the same technological landscape documented by managers.

Figure 8 reveals substantial heterogeneity in worker-reported AI adoption across occupations. Managers are by far the most likely group to report AI use at the workplace, followed by professionals, while clerical, service, skilled traders, and production workers report substantially lower adoption rates. This occupational gradient is consistent with the firm-level evidence that adoption is concentrated in firms with stronger digital capabilities and more skill-intensive workforces, and it also suggests differences in awareness within firms, with managers having greater visibility into firm-wide adoption. This is particularly relevant for interpreting the firm-level survey, where managers are the primary informants.

Figure 8: Worker-Reported Adoption of AI in the Workplace by Occupation

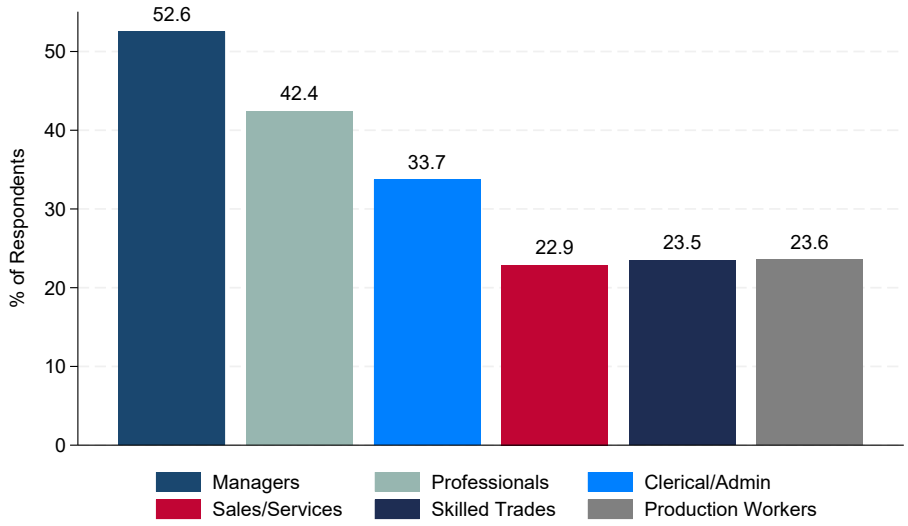

Note: The figure reports the share of individuals who report that the technology is used at their workplace by occupational group. Core AI technologies are: machine learning, natural language processing, voice recognition, and machine vision. Statistics are weighted using survey weights.

4.2 Individual-Level Evidence of AI Adoption

We next examine individual-level patterns of AI adoption using regression analysis. We estimate logit models in which the dependent variable is an indicator for whether an individual reports AI use at the workplace. The regressions are estimated on the subsample of employed respondents, for whom workplace characteristics can be linked to firm-level data from DST registers. The analysis relates worker-reported adoption to individual characteristics, occupation, and firm characteristics, allowing us to assess whether occupational differences persist once observable worker and firm attributes are taken into account, and whether worker-level reports reproduce the firm-level correlates of adoption documented in Section 3.2.

Our baseline specification is the following:

$$\Pr\big(AI_i = 1\big) = \Lambda\big(X_i'\beta + Z_j'\delta + \theta_o + \theta_s + \gamma_r\big) \tag{4}$$

where AI_i is an indicator equal to one if an individual i reports AI use at the current or most recent workplace at firm j. The vector X_i collects individual characteristics, including age group indicators, gender, foreign background, educational attainment, STEM background, and employment status. The vector Z_j includes firm-level characteristics of the individual's most current workplace, such as firm size, value added per worker, ownership structure, and import and export activity. The term θ_o denotes occupation dummies. Finally, θ_s and γ_r capture industry and province fixed effects, respectively.[12]

12. Regressions are weighted by survey weights, and standard errors are robust to heteroskedasticity.

The worker-level adoption outcomes are measured in survey waves conducted in 2023 and 2024, while firm-level and individual characteristics are obtained from administrative registers measured in 2022 and treated as predetermined.

Figure 9 and Table F.7 reports average marginal effects from the worker-level logit models of AI adoption at the workplace. Several interesting patterns emerge.

Figure 9: AI Adoption at the Workplace: Individual-Level Evidence

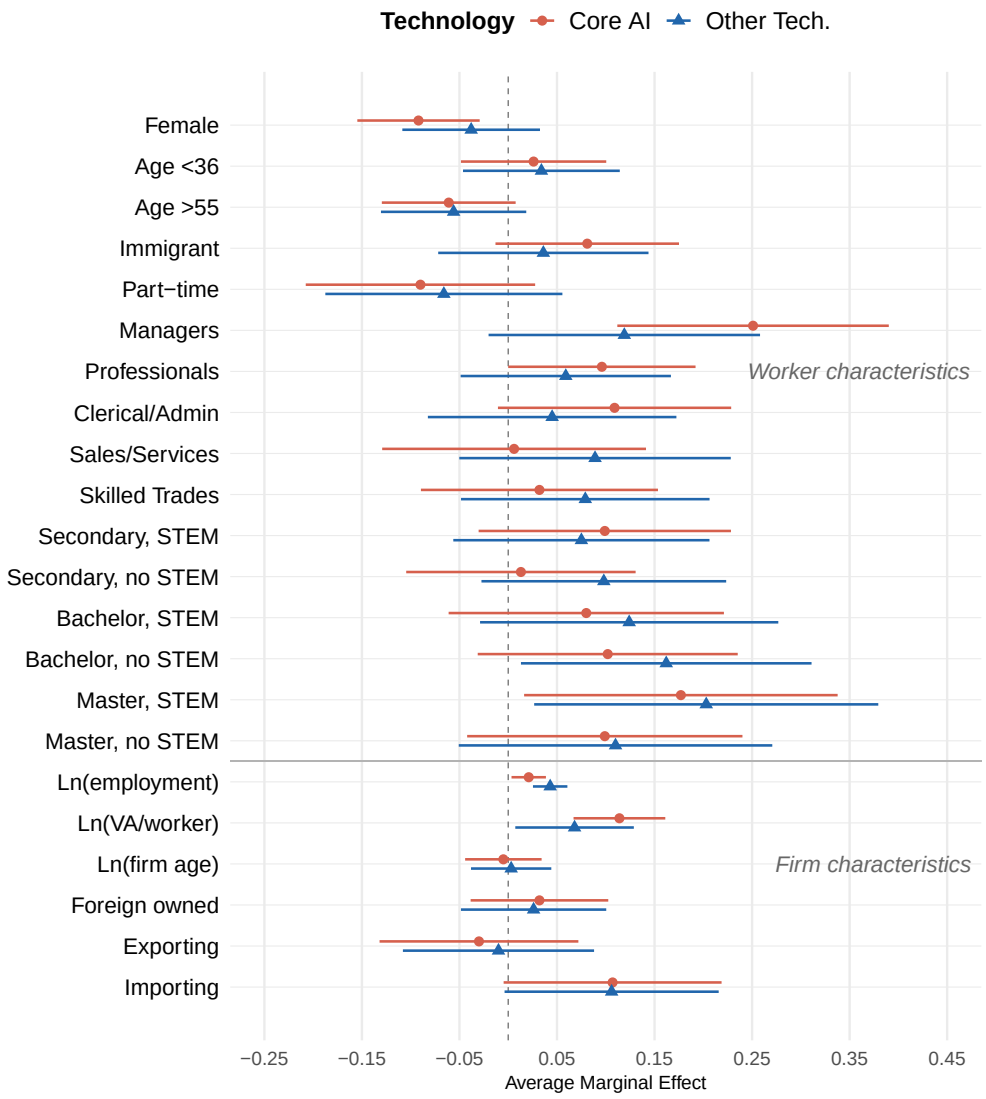

Note: The figure displays average marginal effects (AMEs) from logit models of technology adoption at the extensive margin, estimated separately for two technology categories: Core AI (Machine Learning, Natural Language Processing, Machine Vision, or Voice Recognition Software), Other Technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems). Specifications are estimated on subsample of employed individuals and correspond to columns (2) and (4) in Table F.7. Whiskers denote 95% confidence intervals based on robust standard errors. Worker characteristics are binary indicators, as Exporting, Importing and Foreign owned. Omitted categories are: males, workers aged 36–54, natives, individuals with below-secondary education, and production workers. Regressions are weighted using survey weights.

The regression results confirm a strong occupational gradient in reported AI use, particularly for core AI technologies. Managers are significantly more likely than production workers—the omitted occupational category—to report core AI adoption, even

after controlling for education and other individual attributes and firm characteristics. Professionals also exhibit significantly higher adoption probabilities, while differences for lower-skilled occupations are small or insignificant once controls are included. This pattern mirrors the descriptive evidence and is consistent with managers being more aware of firm-wide adoption of frontier AI technologies, especially when AI is embedded in existing software interfaces or machinery that non-managerial workers use without necessarily recognizing the underlying technology. For other advanced digital technologies, the occupational gradient is weaker and less precisely estimated once firm characteristics are included, suggesting that these technologies are more broadly diffused across tasks and occupations.

Beyond occupation, individual human capital characteristics are important predictors of worker-reported adoption. Tertiary-educated workers are significantly more likely to report AI adoption, with effects that are large in magnitude and robust across specifications. The gradient is present for both STEM and non-STEM education, although somewhat stronger for STEM fields. Younger and foreign workers are also more likely to report AI use, while gender differences are modest once controls are included.

Worker-level regressions also replicate several key firm-level patterns. Workers employed in larger firms and in firms with higher value added per worker are significantly more likely to report adoption. We also find a positive association between importing activity and worker-reported adoption, consistent with the firm-level evidence that internationally engaged firms are more likely to adopt advanced technologies.

Overall, these findings highlight that awareness of AI adoption in the workplace is concentrated among workers in managerial and high-skill roles and varies systematically across occupational groups. This pattern likely reflects both differential exposure to AI deployment and differences in visibility into organizational technology decisions that are closely linked to occupational position and human capital.

4.3 AI Knowledge and Workforce Readiness

We next examine self-assessed AI knowledge as a measure of workforce readiness. While adoption captures whether AI is present in the workplace, knowledge reflects whether workers are equipped to engage with these technologies.

Descriptive evidence shows that self-assessed AI knowledge varies systematically with human capital and occupational position. Workers with tertiary education—particularly those with a STEM background—report substantially higher levels of AI knowledge, while workers with lower education levels are much more likely to report no knowledge (Figure 10). Likewise, managers and professionals report much higher AI knowledge than production and service workers (Figure E.2.8). These patterns support the firm-level

evidence in Section 3.2, where AI adoption is concentrated in firms with more highly educated and skill-intensive workforces.

Figure 10: Self-Assessed Knowledge of AI by Education Level

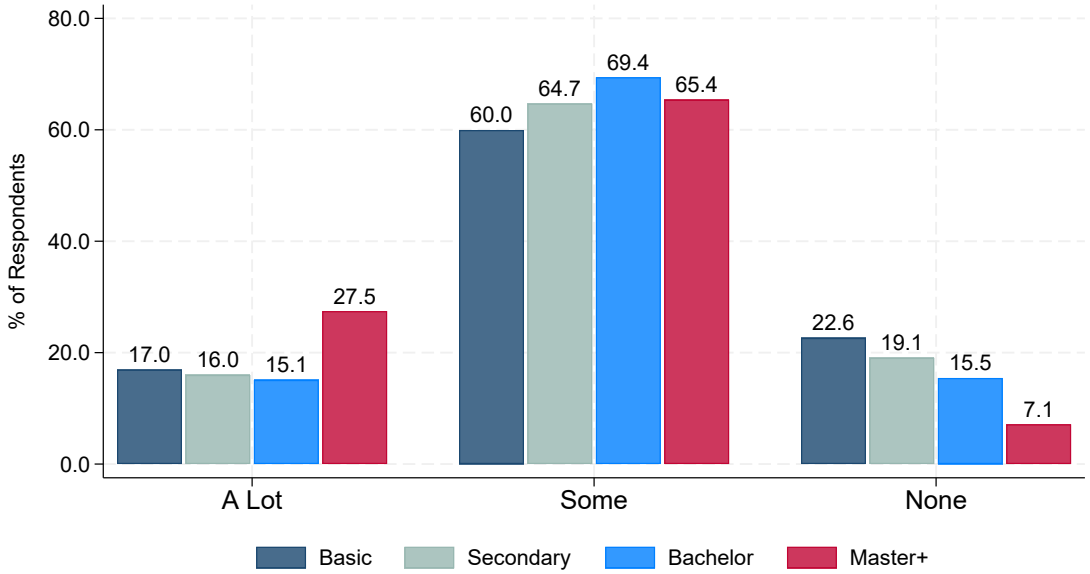

Note: The figure reports self-assessed knowledge of AI by educational level, based on the question: "How much would you say you personally know about artificial intelligence (AI)?" Each bar represents the share of individuals within each education group who report having a given level of AI knowledge shown on the horizontal axis. Statistics are weighted using survey weights.

To formalise these relationships, Table F.8 presents logit regressions of self-assessed AI knowledge. Column (1) uses the full individual sample and includes labour market status indicators, allowing us to compare AI knowledge across employed and non-employed groups, including students and retirees. Columns (2) and (3) restrict the analysis to employed respondents and mirror the specification used in the workplace AI adoption regressions in equation (4), making these columns directly comparable across outcomes. Column (3) additionally includes firm-level characteristics of the respondent's workplace.

The results from the full sample in column (1) reveal clear and systematic gradients in AI knowledge across demographic groups and labour market status. Younger individuals report significantly higher AI knowledge, while older individuals report substantially lower levels. These age patterns are closely linked to life-cycle status: students exhibit particularly high levels of AI knowledge, whereas retired individuals report markedly lower levels. Gender differences are also pronounced, with women reporting lower AI knowledge than men.

Human capital variables emerge as the dominant predictors of AI knowledge, consistent with the firm-level evidence in Section 3.2 on the importance of workforce composition for AI adoption. Relative to individuals with basic education, those with secondary education report higher AI knowledge, while the differences become substantially larger at the tertiary level. These effects are particularly strong for individuals with a STEM

background, but are also clearly present for non-STEM university graduates, indicating that both technical and general higher education are closely associated with AI-related knowledge. Managers, professionals, and clerical workers report significantly higher AI knowledge than production workers, even after controlling for education, age, and other individual characteristics.

In contrast, firm-level characteristics contribute little to explaining variation in self-assessed AI knowledge once individual characteristics and occupations are accounted for. Firm size, productivity, and trade exposure are not systematically associated with reported knowledge. The main exception is foreign ownership, which is positively associated with AI knowledge and likely reflects greater exposure to international technological frontiers in multinational firms. While the inclusion of firm-level variables reduces the estimation sample, the coefficients on individual characteristics remain stable in both magnitude and precision.

To sum up, these results highlight that AI knowledge is primarily shaped by individual human capital and occupational roles rather than firm characteristics. The strong and consistent gradients by education, STEM training, age, and occupation mirror the firm-level findings and reinforce the central role of workforce composition in understanding both the adoption and the use of AI technologies.

5 AI Exposure Measures and Observed AI Adoption

Much of the empirical literature studies the effects of artificial intelligence using AI exposure measures rather than observed adoption decisions. These exposure indices typically map AI-related technological capabilities to occupations based on task and ability descriptions, most commonly drawing on information from O*NET, and are widely used to proxy firms' exposure to AI-driven technological change. While such measures aim to capture heterogeneity in task content and occupational roles—dimensions that are central in our worker-level evidence—their ability to predict actual AI adoption at the firm level has rarely been assessed. In this section, we therefore use our survey-based measures of observed AI adoption to evaluate the predictive power of commonly used AI exposure indices. This exercise provides both a validation of these widely used proxies and a test of their relevance for empirical research and policy.

We estimate a series of linear probability models of the following form:

$$Y_j = \alpha + \beta\, \text{AIExposure}_{j,2015} + X'_{j,2015}\gamma + \theta_s + \delta_p + \varepsilon_j \tag{5}$$

where the dependent variable is a binary indicator for whether a firm in 2023/2024 has adopted any of the four technologies classified as "Core AI" in section 2.2. The specification compares firms within the same sector and province but with different AI exposure levels

in 2015, according to different indices used in the literature.[13] All exposure measures and controls are measured in 2015, well before the adoption decisions observed in 2023/2024, to ensure that occupational composition is predetermined and not itself a consequence of AI adoption. Controls include firm size measured by employment quintiles, firm revenues measured in quintiles, firm age, indicators for foreign ownership, importing status, and exporting status, measures of exposure to software and robot automation based on Webb 2019, and the firm's baseline share of blue-collar workers.

Table 3: AI Exposure and AI Adoption

Exp. Measure	(1) Webb (2019)	(2) Felten et al. (2018)	(3) Felten et al. (2021)	(4) Eloundou et al. (2023)	(5) Engberg et al. (2024)
Dep. Var.: AI adopted (survey)					
AI exposure	0.007	-0.028	0.069*	0.038	0.089***
	(0.020)	(0.018)	(0.038)	(0.026)	(0.032)
Sector FE	✓	✓	✓	✓	✓
Province FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
R-squared	0.208	0.210	0.211	0.210	0.214
Within R^2	0.116	0.118	0.119	0.118	0.123
F-stat	11.30	11.65	12.17	11.59	13.24
F-KP	0.10	2.48	3.26	2.19	7.71
Observations	2,799	2,799	2,799	2,799	2,799

Note: The dependent variable in an indicator for adopting any of the following AI technologies: Machine Learning (ML), Natural Language Processing (NLP), Machine Vision (MV), and Voice Recognition Software (VRS). Each column reports the β from specification (5), estimated on separate models using a different AI exposure measure reported on columns (all standardised). Controls, measured in 2015, include firm size measured by employment quintiles, firm revenues measured in quintiles, firm age, indicators for foreign ownership, importing status, and exporting status, measures of exposure to software and robot automation based on Webb (2019), and the firm's baseline share of blue-collar workers. All regressions are weighted by the product of the estimated inverse probability of responding to the survey and firm size in 2015. Robust standard errors are reported in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3 presents results from equation (5) for five AI exposure measures, all constructed from the occupational composition of firms in 2015. The measures differ in how they capture AI capabilities: Webb 2019 uses patent text linked to occupational task descriptions; Felten et al. 2018 and Felten et al. 2021 use expert-curated benchmarks of AI progress mapped to O*NET abilities, with the latter incorporating more recent and granular benchmarks; Eloundou et al. 2023 focuses specifically on large language model capabilities matched to O*NET tasks;[14] and Engberg et al. 2024 construct a time-varying index (DAIOE) by tracking AI capability benchmarks across subdomains year by year.[15]

13. We use 2015, because it predates the onset of AI diffusion and thus reflects occupational structures not yet shaped by AI adoption.

14. The table reports the version based on human ratings of short-term task exposure. A similar effect is found with the longer-term version.

15. Column 5 uses the occupational composition in 2015 and the variation in occupational exposure between 2015 and 2022.

Since the occupational composition underlying the different exposure indices is the same, the results in Table 3 indicate that the type of information used to capture AI capabilities and the way this information is matched to occupations matter for the ability to predict actual AI adoption as measured in the survey. In particular, exposure measures based on patent data (Webb 2019) and those with a narrow focus on large language models (Eloundou et al. 2023) do not appear to be strong predictors of AI adoption. This suggests that some frontier AI innovations may occur in functionalities that lack immediate business applications and are not readily adoptable by firms.

By contrast, measures that rely on benchmark-based indicators of AI progress, such as those drawing on sources like the Electronic Frontier Foundation (EFF), and that link AI capabilities to occupations through expert assessments of the abilities required to perform job tasks exhibit greater predictive power (Felten et al. 2021, Engberg et al. 2024). In particular, the exposure measure proposed by Engberg et al. 2024, which captures the evolution of AI technological capabilities over time, displays the strongest association with AI adoption among Danish firms. For robustness, Tables F.9 and F.10 in the Appendix show that these exposure indices are not capturing a general propensity of AI-adopting firms to adopt technologies more broadly, as they exhibit no predictive power for robot adoption or for the adoption of other advanced technologies.[16] Having confirmed that exposure does not imply adoption, this section shows that how AI exposure is measured matters. Moreover, it points to the importance of studying the barriers and predictors of AI adoption, on which this paper provides novel evidence.

6 Conclusions

This paper studies AI adoption using linked survey and administrative data from Denmark and moves beyond indirect exposure proxies. Our central finding is that AI adoption is technology-specific: the firm characteristics associated with adoption vary across technology types, and treating AI as a monolithic category obscures meaningful heterogeneity. Core AI and Generative AI have distinct predictor profiles, with STEM-educated workforces associated with core AI adoption and non-STEM university-educated workforces predicting Generative AI adoption. Within adopting firms, machine learning and natural language processing are associated with use across multiple business functions, consistent with the broader applicability theorised for general-purpose technologies, while other advanced technologies remain function-specific. Moreover, the capabilities associated with adoption differ from those predicting the breadth of deployment: workforce composition drives the extensive margin, while organisational scale and digital maturity

16. Other advanced technologies are: Automated Guided Vehicles (AGV), Automated Storage Systems (ASS), Radio Frequency Identification (RFID), Touch Screens/Kiosks (TS), Robotics, Augmented Reality (AR), and Automated Decision-Making Systems (ADMS).

matter more at the intensive margin. Individual-level evidence provides a consistent micro-level counterpart: AI use and knowledge are systematically higher among younger and more educated workers, particularly in high-skill occupations, reinforcing that workforce composition is central to both the adoption and effective use of AI technologies. Finally, commonly used AI exposure indices vary substantially in their ability to predict which firms actually adopt, with benchmark-based capability measures outperforming patent-based and LLM-focused alternatives. The divergence between exposure and adoption underscores that studying AI's economic effects through indirect proxies alone may miss the organisational factors associated with which firms actually adopt these technologies.

The patterns emerging from our analysis suggest that AI adoption currently reinforces existing firm heterogeneity rather than acting as a uniform technological shock. Firms that are already digitally mature and skill-intensive adopt earlier and deploy these technologies more broadly, whereas firms lacking these complementary capabilities remain at the margin. Because different technologies are associated with different complementary inputs, the distributional consequences of AI will depend on which technologies diffuse and across which firms.

A key direction for future research is to examine how the patterns documented here extend across institutional and economic contexts. This study focuses on Denmark, a small, high-income and highly digitalised economy, which provides a setting where AI adoption is already sufficiently widespread to allow for a detailed analysis of both adoption and deployment patterns within firms. As a result, the findings offer a useful benchmark for understanding how different technologies diffuse and are integrated in advanced economies. Extending this analysis to countries at earlier stages of digitalisation, or with different institutional environments, would provide valuable insights into how the determinants of adoption and the organisation of technology use vary across settings. In this sense, the Danish context may be particularly informative as an early case, highlighting patterns that may emerge more broadly as AI adoption continues to expand across Europe and beyond.

The findings have important policy implications. If policymakers seek to influence AI adoption, our results suggest that attention to complementary capabilities – digital infrastructure, data management, and workforce skills – may matter more than access to AI technologies alone. Because core AI and generative AI are associated with different workforce characteristics, effective training and education policies may need to be tailored to the technology landscape firms face, rather than treating AI skills as homogeneous. More broadly, the results underscore that the economic consequences of AI will be shaped not only by technological progress, but by how heterogeneous technologies diffuse across firms and interact with complementary capabilities.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, tasks and technologies: Implications for employment and earnings." In *Handbook of labor economics,* 4:1043–1171. Elsevier.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and jobs: Evidence from US labor markets." *Journal of political economy* 128 (6): 2188–2244.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–2451. https://doi.org/10.3982/ECTA13408. https://www.jstor.org/stable/43866416.

Aghion, Philippe, Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard. 2025. "How Different Uses of AI Shape Labor Demand: Evidence from France." *AEA Papers and Proceedings* 115 (May): 62–67. https://doi.org/10.1257/pandp.20251047. https://www.aeaweb.org/articles?id=10.1257/pandp.20251047.

Alesina, Alberto, and Eliana La Ferrara. 2005. "Ethnic Diversity and Economic Performance." *Journal of Economic Literature* 43 (3): 762–800.

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

Bloom, Nicholas, Mirko Draca, and John Van Reenen. 2016. "Trade induced technical change? The impact of Chinese imports on innovation, IT and productivity." *The review of economic studies* 83 (1): 87–117.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "The organization of firms across countries." *The quarterly journal of economics* 127 (4): 1663–1705.

Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and explaining management practices across firms and countries." *The quarterly journal of Economics* 122 (4): 1351–1408.

Bonney, Kathryn, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, and Keith Savage. 2024. *Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey.* Working Paper, Working Paper Series 32319. National Bureau of Economic Research, April. https://doi.org/10.3386/w32319. http://www.nber.org/papers/w32319.

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

Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond computation: Information technology, organizational transformation and business performance." *Journal of Economic perspectives* 14 (4): 23–48.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics* 13, no. 1 (January): 333–72. https://doi.org/10.1257/mac.20180386. https://www.aeaweb.org/articles?id=10.1257/mac.20180386.

Bryson, Alex, and John Forth. 2023. "Worker Representation." In *Encyclopedia of Labour Studies,* edited by Tor Eriksson, 235–239. Cheltenham, UK: Edward Elgar. ISBN: 978-1800377530.

Calvino, Flavio, and Luca Fontanelli. 2023. *A portrait of AI adopters across countries: Firm characteristics, assets' complementarities and productivity.* Technical report, OECD Science, Technology and Industry Working Papers 2023/02. Paris: OECD Publishing. https://doi.org/10.1787/0fb79bb9-en. https://dx.doi.org/10.1787/0fb79bb9-en.

Cohen, Wesley M, Daniel A Levinthal, et al. 1990. "Absorptive capacity: A new perspective on learning and innovation." *Administrative science quarterly* 35 (1): 128–152.

Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible capital and US economic growth." *Review of income and wealth* 55 (3): 661–685.

Deming, David J. 2017. "The growing importance of social skills in the labor market." *The Quarterly Journal of Economics* 132 (4): 1593–1640.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "Gpts are gpts: An early look at the labor market impact potential of large language models." *arXiv preprint arXiv:2303.10130* 10.

Engberg, Erik, Holger Görg, Magnus Lodefalk, Farrukh Javed, Martin Längkvist, Natália Pimenta Monteiro, Hildegunn Kyvik Nordås, Giuseppe Pulito, Sarah Schroeder, and Aili Tang. 2024. *AI unboxed and jobs: A novel measure and firm-level evidence from three countries.* Technical report. RF Berlin-CReAM Discussion Paper Series.

Eurostat. 2026. "Artificial intelligence by size class of enterprise." Dataset. Accessed March 26, 2026. https://doi.org/10.2908/ISOC_EB_AI. https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ai/default/table?lang=en.

Felten, Edward, Manav Raj, and Robert Seamans. 2018. "A Method to Link Advances in Artificial Intelligence to Occupational Abilities." *AEA Papers and Proceedings* 108 (May): 54–57. https://doi.org/10.1257/pandp.20181021. https://www.aeaweb.org/articles?id=10.1257/pandp.20181021.

———. 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic Management Journal* 42 (12): 2195–2217.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2010. "Regularization Paths for Generalized Linear Models via Coordinate Descent." *Journal of Statistical Software* 33 (1): 1–22. https://doi.org/10.18637/jss.v033.i01.

Garicano, Luis, and Esteban Rossi-Hansberg. 2006. "Organization and inequality in a knowledge economy." *The Quarterly journal of economics* 121 (4): 1383–1435.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital economics." *Journal of economic literature* 57 (1): 3–43.

Graetz, Georg, and Guy Michaels. 2018. "Robots at work." *Review of economics and statistics* 100 (5): 753–768.

Hong, Lu, and Scott E. Page. 2004. "Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers." *Proceedings of the National Academy of Sciences* 101 (46): 16385–16389. https://doi.org/10.1073/pnas.0403723101. https://doi.org/10.1073/pnas.0403723101.

Humlum, Anders, and Emilie Vestergaard. 2024. *The adoption of ChatGPT*. Technical report. IZA Discussion Papers.

Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti. 2008. "Skill Dispersion and Firm Productivity: An Analysis with Employer–Employee Matched Data." *Journal of Labor Economics* 26 (2): 247–285. https://doi.org/10.1086/587430. https://doi.org/10.1086/587430.

Keller, Wolfgang. 2004. "International technology diffusion." *Journal of economic literature* 42 (3): 752–782.

Kostøl, Fredrik B., and Elin Svarstad. 2023. "Trade Unions and the Process of Technological Change." *Labour Economics* 84:102386.

Lee, Yong Suk, Taekyun Kim, Sukwoong Choi, and Wonjoon Kim. 2022. "When does AI pay off? AI-adoption intensity, complementary investments, and R&D strategy." *Technovation* 118:102590.

Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *The Review of Economic Studies* 70 (2): 317–341. https://doi.org/10.1111/1467-937X.00246. https://academic.oup.com/restud/article-abstract/70/2/317/1586773.

McElheran, Kristina, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. 2024. "AI adoption in America: Who, what, and where." *Journal of Economics & Management Strategy* 33 (2): 375–415.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. "The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve(s)." Working paper, presented at ASSA 2026.

Navon, Guy. 2009. "Human Capital Spillovers in the Workplace: Labor Diversity and Productivity." *Labour Economics* 16 (1): 109–121. https://doi.org/10.1016/j.labeco.2008.06.001. https://doi.org/10.1016/j.labeco.2008.06.001.

OECD. 2019. *Artificial Intelligence in Society.* Paris: OECD Publishing. https://doi.org/10.1787/eedfee77-en.

Østergaard, Christian R, Bram Timmermans, and Kari Kristinsson. 2011. "Does a different view create something new? The effect of employee diversity on innovation." *Research policy* 40 (3): 500–509.

Parrotta, Pierpaolo, Dario Pozzoli, and Mariola Pytlikova. 2014a. "Labor Diversity and Firm Productivity." *European Economic Review* 66:144–179. https://doi.org/10.1016/j.euroecorev.2013.12.002. https://doi.org/10.1016/j.euroecorev.2013.12.002.

———. 2014b. "The Nexus between Labor Diversity and Firms' Innovation." *Journal of Population Economics* 27 (2): 303–364. https://doi.org/10.1007/s00148-013-0481-8. https://doi.org/10.1007/s00148-013-0481-8.

Rammer, Christian, Gastón P Fernández, and Dirk Czarnitzki. 2022. "Artificial intelligence and industrial innovation: Evidence from German firm-level data." *Research Policy* 51 (7): 104555.

Tauman, Yair, and Yoram Weiss. 1987. "Labor Unions and the Adoption of New Technology." *Journal of Labor Economics* 5 (4): 477–501.

Webb, Michael. 2019. "The Impact of Artificial Intelligence on the Labor Market." *SSRN Electronic Journal* (January). https://doi.org/10.2139/ssrn.3482150.

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J Davis, Kevin M Foster, Aaron Jalca, et al. 2026. *Firm Data on AI.* Working Paper, Working Paper Series 34836. National Bureau of Economic Research, February. https://doi.org/10.3386/w34836. http://www.nber.org/papers/w34836.

Appendices

Contents

A Survey and Data Preparation **42**
A.1 Employer and Employee Survey – Questionnaires and Implementation . . 42
A.2 Definition of Artificial Intelligence and Related Technologies 42

B Co-adoption Patterns Across Technologies **45**

C Probability of Responding to the Survey **46**

D Lasso Logit Model of AI Adoption **47**

E Descriptives **52**
E.1 Firm-level Survey: AI Adoption Patterns 52
E.2 Individual-Level Survey: Additional Figures 55

F Main Results - Tables & Figures **56**

A Survey and Data Preparation

A.1 Employer and Employee Survey – Questionnaires and Implementation

This section documents the employer survey questionnaire used in our study of AI adoption.

The questionnaire follows a common structure for all firms, with items grouped into blocks covering industrial relations, digitalisation, adoption and use of AI technologies, sourcing and employee capabilities, randomized information treatments, expectations for future adoption, perceived impacts, and views on governance.

The survey was administered online (Web CAWI) and hosted by Statistics Denmark, which also managed the invitations. On the introduction screen, respondents were informed that the survey was conducted on behalf of a research group at Aarhus University, accompanied by the standard Statistics Denmark confidentiality and data-protection notice, and that the survey concerns "technology, employment, policy, and society". Respondents answered on behalf of their firm, not in a personal capacity.

For brevity, the description below focuses on the core AI modules—technology checklist and application domains (Block C), sourcing and capabilities (Block D), the short randomized information texts shown immediately before the future-use items (Block E→F), and expectations about future adoption (Block F). The full questionnaire—including industrial relations (A), digitalization baseline (B), perceived impacts on staffing and task content as well as employee well-being (G), and views on governance and public spending (H)—is available in our replication package (see Appendix Survey Questionnaire). In the Online Appendix, we present the English version of the questionnaire.[17]

A.2 Definition of Artificial Intelligence and Related Technologies

Definition of Artificial Intelligence

We follow the OECD (2019) definition of artificial intelligence as "a machine-based system that can influence the environment by producing an output (predictions, recommendations, or decisions) for a set of objectives." AI systems use machine- or human-generated inputs to perceive environments, structure these inputs into models, and generate inferences that inform possible actions. This operational definition aligns with the technologies covered in the Danish Digitilisation and AI Survey.

The survey distinguishes between *Core AI*, *Generative AI*, and a broader set of *Other Advanced Technologies* that either complement AI or incorporate AI-enabled components. Table A.2.1 summarises the operational definitions used in the questionnaire.

17. The original Danish version of the questionnaire is available from the authors upon request.

Table A.2.1: Definitions of AI and Related Technologies Used in the Survey

Technology	Definition
Machine Learning	Algorithms that learn patterns from data to improve predictions or performance without explicit reprogramming.
Machine Vision	Systems enabling automated inspection, recognition, or interpretation of images or video, typically used for quality control or object detection.
Natural Language Processing	Techniques that allow computers to process, interpret, and generate human language in written or spoken form.
Voice Recognition	Software that converts speech into text or interprets spoken commands, often combined with NLP for downstream applications.
Generative AI	AI systems that produce new content (such as text, images, audio, or code) by learning patterns from large training datasets, including large language models.
Automated Guided Vehicles	Computer-controlled vehicles that move autonomously using sensors and software, without a human driver.
Automated Storage and Retrieval Systems	Systems designed to automatically locate, retrieve, and return items from predefined storage locations.
Radio-Frequency Identification	Technologies using tags and radio-wave readers to identify and track objects, either through fixed scanners or mobile devices.
Robotics	Reprogrammable machines capable of executing complex or repetitive physical tasks autonomously or semi-autonomously.
Touchscreens/Kiosks	Customer-facing digital interfaces that facilitate information access or transactions, such as self-checkout, ordering, or check-in terminals.
Augmented Reality	Systems that overlay computer-generated information or graphics onto the user's view of the real world.
Automated Decision Making Systems	Software systems that apply algorithmic rules or models to make decisions without human intervention (e.g., credit scoring, fraud detection, risk classification).

Notes: The table reports descriptions of the technologies included in the survey. Technologies in the top panel are classified as *Core AI*, while those in the bottom panel are classified as *Other Advanced Technologies*.

Digital Infrastructure, AI Technologies, and Business Functions

In addition to technology-specific information, the survey provides rich measures of firms' digital readiness, which are strongly associated with AI adoption. First, firms evaluate the importance of cloud infrastructure through the question: *"How important are cloud services in general to the company's IT functions?"*, with response categories 1 = not important at all, 2 = somewhat important, 3 = pretty important, and 4 = very important.[18] This ordered measure provides a direct indicator of cloud readiness, reflecting the availability of flexible computing resources and interoperable data systems that lower the effective costs of deploying AI applications.

Second, we construct a *Digitalisation Index* capturing the extent to which key categories of business information were stored in digital format in 2022. Firms report, on an ordered scale from 1 ("none") to 4 ("all of it"), how digitalised six domains are: Human Resources, Finance, Customer Feedback, Marketing, Supply Chain, and Production. Summing the responses yields an index ranging from 0 to 18, where higher values indicate more pervasive digitalisation of the firm's operations. This measure captures the foundational data infrastructure that underpins AI implementation, and it serves as a proxy for the firm's broader technological absorptive capacity. Both cloud readiness and digitalisation strongly correlate with observed adoption patterns documented in the descriptive and empirical analysis that follows.

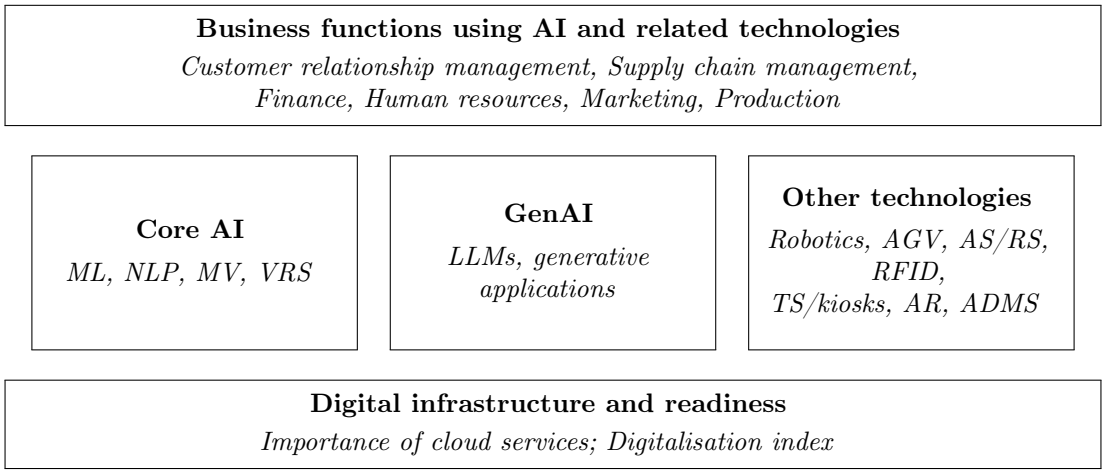

Figure A.2.1: Digital infrastructure, AI technologies, and business functions

Figure A.2.1 provides a conceptual taxonomy linking the technologies, business functions, and digital capabilities observed in the Danish Digitisation and AI Survey. At the centre of the figure are three groups of technologies. These technologies are used across a wide set of business functions that structure firms' internal and external activities: customer relationship management, supply chain management, finance, human resources, marketing, and production. The survey records whether AI or other technologies are used

18. Responses 98 = hidden and 99 = "don't know" are treated as missing

within each of these functional areas, allowing us to analyse heterogeneity in adoption across organisational domains.

At the base of the structure are capabilities associated with a firm's readiness to deploy and scale AI: the reported importance of cloud services for the firm's IT operations and the composite digitalisation index. These variables proxy for the quality of a firm's data infrastructure and its capacity to integrate AI into existing workflows. Taken together, the figure clarifies how the empirical components in our analysis fit into a coherent structure: digital infrastructure is associated with technology adoption; technologies differ in their complexity and complementarities; and their use varies across business functions in ways that reflect organisational needs and capabilities.

B Co-adoption Patterns Across Technologies

To explore complementarities and co-occurrence patterns among technologies, we compute a Jaccard similarity matrix across all pairs of technologies. For any two technologies A and B, the Jaccard index is defined as

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|},$$

which captures the share of adopters using both technologies relative to the share using at least one of them. Figure B.1 presents the resulting heatmap. Higher values indicate technologies that tend to be adopted jointly, while lower values reflect technologies that are more standalone or specialised. These patterns provide useful descriptive evidence on the structure of technology portfolios within firms.

Figure B.1: Jaccard Co-adoption Matrix Across Technologies

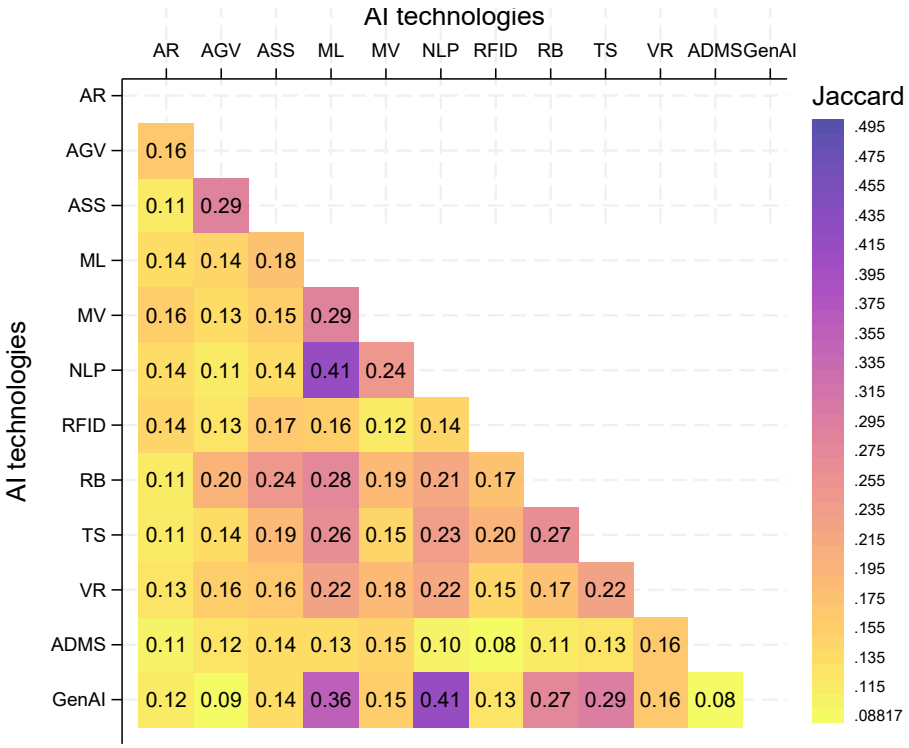

Note: The figure reports the Jaccard similarity index for all pairs of technologies, defined as $J(A,B) = \frac{|A \cap B|}{|A \cup B|}$, where A, B are technologies. Higher values indicate that two technologies are frequently adopted together by the same firms. The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Other advanced technologies are AR: Augmented Reality, AGV: Automated Guided Vehicles, ASS: Automated Storage Systems, RFID: Radio Frequency Identification, Robotics, TS: Touch Screens/Kiosks, ADMS: Automated Decision-Making Systems. Sample includes all firms in the 2024 survey waves, which includes information on GenAI.

C Probability of Responding to the Survey

The sample was stratified by Statistics Denmark (DST) across 18 industry groups and five size categories and comprises 10,083 firms that are representative of the population at the start of data collection. Nevertheless, survey participation and completion may be endogenous with respect to firm characteristics and correlated with the propensity to adopt AI. Survey weights provided by DST and based on the stratification design can partially address these concerns. To address selection into survey participation beyond the stratification variables and to flexibly adjust for observable determinants of survey response, we leverage rich administrative register data to estimate firms' response probabilities and construct inverse probability weights.

We estimate each firm's probability of answering the AI–adoption survey using a logistic lasso on the universe of Danish firms with at least five employees. Let $R_j \in \{0,1\}$ indicate whether firm j responded, X_j be the vector of covariates,[19] and let δ_s and γ_f

19. The list of covariates includes: firm age; firm age squared; dummies for province, indicator for foreign ownership; share of female employees; share of immigrant employees; share of employees younger

denote, respectively, industry and firm-size fixed effects that coincide with the DST stratification used in the survey design. These fixed effects are always included and not penalized.

The fitted probability is

$$\hat{p}_j = \Lambda\Big(X_j'\hat{\beta} + \delta_s + \gamma_f\Big), \qquad \Lambda(t) = \frac{1}{1+e^{-t}},$$

where $(\hat{\delta}, \hat{\gamma}, \hat{\beta})$ solves the lasso–penalized log-likelihood

$$(\hat{\delta}, \hat{\gamma}, \hat{\beta}) = \arg\min_{\alpha,\delta,\gamma,\beta} \left\{ -\frac{1}{N}\sum_{i=1}^{N}\Big[R_j \log\Lambda(\eta_j) + (1-R_j)\log\big(1-\Lambda(\eta_j)\big)\Big] + \lambda\sum_{k\in K}|\beta_k| \right\},$$

with linear index $\eta_j = X_j'\beta + \delta_s + \gamma_f$ and where the ℓ_1 penalty is applied only to coefficients $\{\beta_k : k \in K\}$ of the non–fixed-effect covariates.

This approach performs automatic variable selection and shrinkage, mitigating overfitting and multicollinearity and improves predictions and out-of-sample calibration, versus an unpenalized logit when some predictors are weak or collinear.[20] For all regression analyses and descriptive statistics based on survey respondents, we apply inverse probability weights $w_j = \frac{1}{\hat{p}_j}$, constructed from the estimated response probabilities described above. The table of descriptive statistics 1 shows that reweighting by w_j substantially increases the similarity between the respondent sample and the population.

D Lasso Logit Model of AI Adoption

The approach followed in Section 3.1.1 to identify the main predictors of AI adoption relies on a predetermined set of covariates drawn from the literature. In this section, we implement an alternative, more data-driven empirical strategy in which predictors of AI adoption are selected from a larger pool of covariates based on their predictive power. An advantage of this framework is that variable selection is not sensitive to an arbitrary ordering of covariates, as predictors are selected based on their stable contribution to predictive performance rather than on sequential inclusion.

than 36; share of employees aged 55 or older; log total exports; log total imports; log value added; log physical capital stock (million DKK); log hourly wage; log yearly wage; share of managers; share of blue-collar workers; share of employees with a bachelor's degree or higher; share of employees with a STEM university education; growth of employment; growth of value-added per worker (defined as DHS growth rates between the average over 2018-2021 and 2022, the last year available).

20. The lasso logit procedure implements the coordinate descent algorithm by Friedman et al. 2010. The algorithm chooses the optimal λ through a cross-validation approach that minimizes the mean deviance across 10 different folds for each λ on the grid. If the unpenalized logit is adequate, the algorithm will choose a very small λ, essentially returning the unpenalized model.

We adopt a two–step procedure. First, we screen covariates with a lasso logistic regression tuned by 10–fold cross–validation on the full set of register–based variables and survey measures (e.g., degree of digitalisation), selecting the subset that minimizes out–of–sample deviance for the probability of adopting AI. Second, we refit an unpenalised logit on the selected covariates to obtain average marginal effects and heteroskedasticity–robust (Huber–White) standard errors.[21]

The second-step logit model weights observations by the inverse of the estimated probability of responding to the survey, $w_j = 1/\hat{p}_j$, as computed in section C. Specifically we estimate the following model:

$$\Pr(AI_j = 1) = \Lambda\Big(X_j'\hat{\beta} + \delta_s + \gamma_p\Big), \qquad \Lambda(t) = \frac{1}{1+e^{-t}}$$

where AI_j is an indicator for AI adoption in firm j, $\Lambda(\cdot)$ is the logistic function, X_j are firm-level covariates,[22] δ_s, γ_p denote industry and province fixed effects, respectively. We estimate separate models for (i) adoption of core AI technologies, (ii) adoption of generative AI, and (iii) adoption of other technologies.

To examine heterogeneity in the predictors of AI adoption, we also apply the same model separately by broad sector (manufacturing, knowledge-intensive services, trade and transport, and other services). Table D.1 reports results from the lasso–logit model with 10-fold cross-validation, while Table D.2 presents results from the parsimonious lasso–logit specification. Firm size, the degree of digitalisation, and the share of the workforce with a university degree emerge as the strongest predictors of AI adoption. The level of digitalisation is also positively associated with the adoption of generative AI and other AI technologies. In contrast to the adoption of core AI technologies, firms' international trade activity, as captured by import intensity, is positively associated with the adoption of generative AI and other technologies (columns 2 and 3 in Table D.1).

21. The first-step lasso logit is run in two versions, both variants fitted by a coordinate-descent algorithm (Friedman et al. 2010): (i) 10-fold cross-validation selects the penalty parameter λ that minimizes the mean deviance on the held-out folds; (ii) cross-validation with the "one-standard-error" rule returns a more parsimonious set of covariates, choosing the largest λ whose CV deviance lies within one standard error of the minimum, yielding a sparser model with predictive performance statistically indistinguishable from the best. Comparing the two models provides a hint as to which covariates are the strongest predictors of AI adoption.

22. The starting list of covariates in the first step is the following: log employment (full-time equivalent); log value added per worker; log revenues; log firm age; foreign ownership; exporter indicator; importer indicator; log capital intensity; log intangible assets per worker; degree of digitalisation (standardised); cloud importance indicator; share of employees with a STEM secondary education; share of employees with a STEM university education; share of employees with a non-STEM university education; share of employees younger than 36; share of employees aged 55 or older; share of managers; share of white-collar workers; share of female employees; share of immigrant employees; log average hourly wage; unionisation share; growth of employment; growth of value-added per worker (both defined as DHS growth rates between the average over 2018-2021 and 2022, the last year available).

Table D.1: Predictors of AI Adoption - Lasso Logit Model

	(1) **Core AI**	(2) **Other Tech.**	(3) **Gen. AI**
Ln(Employment)	0.032*	0.073***	0.070***
	(0.017)	(0.020)	(0.014)
Ln(Revenue)	0.017	-0.005	
	(0.013)	(0.017)	
Foreign owned	0.005	0.033	
	(0.028)	(0.034)	
Employer association	-0.018	-0.043	
	(0.024)	(0.028)	
Digitalisation	0.079***	0.062***	0.101***
	(0.012)	(0.013)	(0.015)
Cloud importance	0.073**	0.040	0.153***
	(0.032)	(0.033)	(0.037)
No University, STEM	-0.158**	-0.063	
	(0.080)	(0.081)	
University, no STEM	0.163**		
	(0.072)		
University, STEM	0.213***		
	(0.074)		
Age < 36	0.180***		
	(0.058)		
White collar share	0.039		
	(0.053)		
Growth Empl. t-1	0.044	0.039	
	(0.029)	(0.034)	
Growth Value added t-1	-0.011**		
	(0.005)		
Exporter dummy		0.009	
		(0.031)	
Importer dummy		0.040	
		(0.032)	
Ln(Capital inten.)		0.005*	
		(0.003)	
Managers share		-0.046	
		(0.140)	
Female share		0.042	
		(0.068)	
Ln(Hourly wage)		0.077	
		(0.060)	
Sector f.e.	✓	✓	✓
Province f.e.	✓	✓	✓
Observations	2899	2899	1751

Notes: The table reports the marginal effects from logit models of probability of adoption by type of technology. In the first step the covariates are selected by running a Lasso logit model on the full set of firm-level variables (sector and province fixed effects are always included). In the second step a logit model is estimated on the selected covariates to obtain standard errors. Estimates are weighted by the inverse probability of responding to the survey. Generative AI adoption is available only in the second wave of the survey (2024). The initial set of covariates in the Lasso step is the following: log employment (full-time equivalent); log value added per worker; log revenues; log firm age; foreign ownership; exporter indicator; importer indicator; log capital intensity; log intangible assets per worker; degree of digitalisation (standardised); cloud importance indicator; share of employees with a STEM secondary education; share of employees with a STEM university education; share of employees with a non-STEM university education; share of employees younger than 36; share of employees aged 55 or older; share of managers; share of white-collar workers; share of female employees; share of immigrant employees; log average hourly wage; unionisation share; employer association dummy; growth of employment; growth of value-added per worker (both defined as DHS growth rates between the average over 2018-2021 and 2022, the last year available). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.3 presents sector-specific predictors of core AI adoption based on the lasso–logit model. Across all sectors, digitalisation is a strong and consistently positive predictor, while import intensity is positively associated with AI adoption primarily in trade-related sectors. In addition, workforce composition matters: a higher share of employees with a university degree is strongly associated with AI adoption in knowledge-intensive business services, whereas STEM education exhibits a negative association in this sector, likely driven by non-university STEM qualifications.

Table D.2: Predictors of AI Adoption - Lasso Logit Parsimonious Model

	(1)	(2)	(3)
	Core AI	**Other Tech.**	**Gen. AI**
Ln(Employment)	0.045***	0.089***	
	(0.014)	(0.009)	
Ln(Revenue)	0.010		
	(0.011)		
Digitalisation	0.086***	0.071***	0.126***
	(0.012)	(0.011)	(0.013)
Cloud importance	0.099***		
	(0.031)		
Sector f.e.	✓	✓	✓
Province f.e.	✓	✓	✓
Observations	3257	3670	1837

Notes: The table reports the marginal effects from logit models of probability of adoption by type of technology. In the first step the covariates are selected by running a Lasso logit model with the "one-standard-error" rule that returns a more parsimonious set of predictors (sector and province fixed effects are always included). In the second step a logit model is estimated on the selected covariates to obtain standard errors. Estimates are weighted by the inverse probability of responding to the survey. Generative AI adoption is available only in the second wave of the survey (2024). The initial set of covariates in the Lasso step is the following: log employment (full-time equivalent); log value added per worker; log revenues; log firm age; foreign ownership; exporter indicator; importer indicator; log capital intensity; log intangible assets per worker; degree of digitalisation (standardised); cloud importance indicator; share of employees with a STEM secondary education; share of employees with a STEM university education; share of employees with a non-STEM university education; share of employees younger than 36; share of employees aged 55 or older; share of managers; share of white-collar workers; share of female employees; share of immigrant employees; log average hourly wage; unionisation share; employer association dummy; growth of employment; growth of value-added per worker (both defined as DHS growth rates between the average over 2018-2021 and 2022, the last year available). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.3: Lasso Logit Predictors of Core AI Adoption, by Sector

Dep. Var.: Core AI adop.	(1) **Manuf.**	(2) **Trade**	(3) **KIBS**	(4) **Other**
Ln(Employment)	0.025	0.014	0.089***	
	(0.036)	(0.029)	(0.020)	
Ln(Revenue)	0.011	0.064***		0.019
	(0.028)	(0.022)		(0.016)
Digitalisation	0.098***	0.025	0.098***	0.047*
	(0.019)	(0.022)	(0.023)	(0.029)
Cloud importance	0.054	0.101*	0.152**	0.122*
	(0.048)	(0.053)	(0.074)	(0.072)
No University, STEM	-0.177**		-0.458*	-0.317
	(0.084)		(0.272)	(0.231)
University, STEM	0.434***	0.309**		
	(0.140)	(0.153)		
Females	0.124		-0.397***	
	(0.110)		(0.112)	
Ln(Value Added)		0.036	0.006	
		(0.058)	(0.039)	
Foreign owned		0.023		
		(0.042)		
University, no STEM		0.133	0.191*	
		(0.124)	(0.105)	
Age < 36		0.132	0.198	
		(0.163)	(0.152)	
Age ≥ 55		-0.048	-0.021	
		(0.216)	(0.184)	
Immigrants		0.414***	0.060	
		(0.134)	(0.152)	
Growth Value added t-1		0.025		
		(0.021)		
Exporter dummy			0.048	0.023
			(0.049)	(0.068)
Importer dummy			0.011	
			(0.051)	
White collar share			0.181	0.110
			(0.153)	(0.084)
Unionisation			-0.127	
			(0.118)	
Ln(Hourly wage)			-0.008	
			(0.107)	
Growth Empl. t-1			0.086	
			(0.054)	
Employer association				-0.058
				(0.057)
Observations	987	745	725	526

Notes: The table reports the marginal effects from logit models of probability of adoption of Core AI technology. In the first step the covariates are selected by running a Lasso logit model on the full set of firm-level variables. In the second step a logit model is estimated on the selected covariates to obtain standard errors. Estimates are weighted by the inverse probability of responding to the survey. The initial set of covariates in the Lasso step is the following: log employment (full-time equivalent); log value added per worker; log revenues; log firm age; foreign ownership; exporter indicator; importer indicator; log capital intensity; log intangible assets per worker; degree of digitalisation (standardised); cloud importance indicator; share of employees with a STEM secondary education; share of employees with a STEM university education; share of employees with a non-STEM university education; share of employees younger than 36; share of employees aged 55 or older; share of managers; share of white-collar workers; share of female employees; share of immigrant employees; log average hourly wage; unionisation share; employer association dummy; growth of employment; growth of value-added per worker (both defined as DHS growth rates between the average over 2018-2021 and 2022, the last year available). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E Descriptives

E.1 Firm-level Survey: AI Adoption Patterns

Figure E.1.1: Adoption Rates Across Technologies

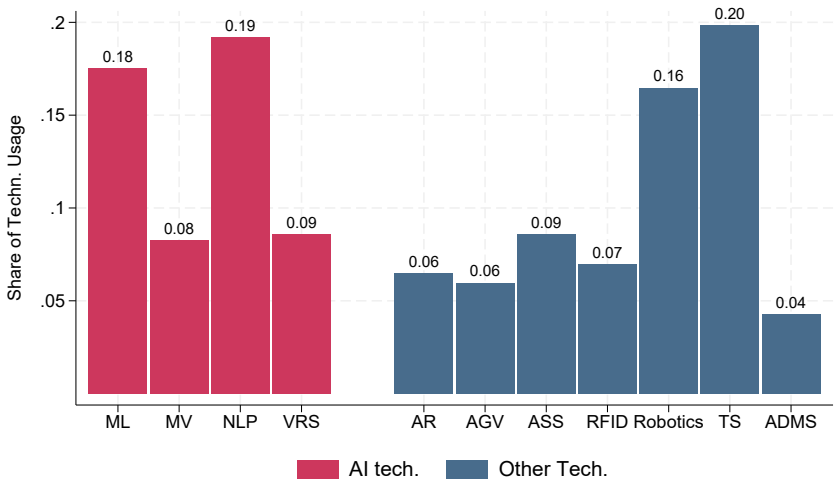

Note: The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Other advanced technologies are AR: Augmented Reality, AGV: Automated Guided Vehicles, ASS: Automated Storage Systems, RFID: Radio Frequency Identification, Robotics, TS: Touch Screens/Kiosks, ADMS: Automated Decision-Making Systems. Generative AI is included only in the 2024 wave of the survey and shows an adoption rate of 50% (not reported in the graph). Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Figure E.1.2: Adoption of AI Technologies vs Other Advanced Technologies by Region

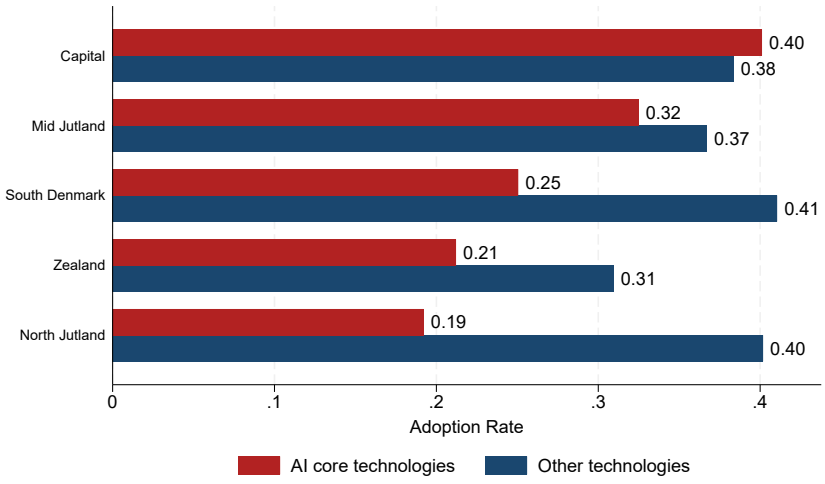

Note: The figure illustrates the weighted adoption rate of core AI and of other advanced technologies by region. The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Other advanced technologies are AR: Augmented Reality, AGV: Automated Guided Vehicles, ASS: Automated Storage Systems, RFID: Radio Frequency Identification, Robotics, TS: Touch Screens/Kiosks, ADMS: Automated Decision-Making Systems. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Figure E.1.3: Internal vs External Development of AI Solutions among Adopters

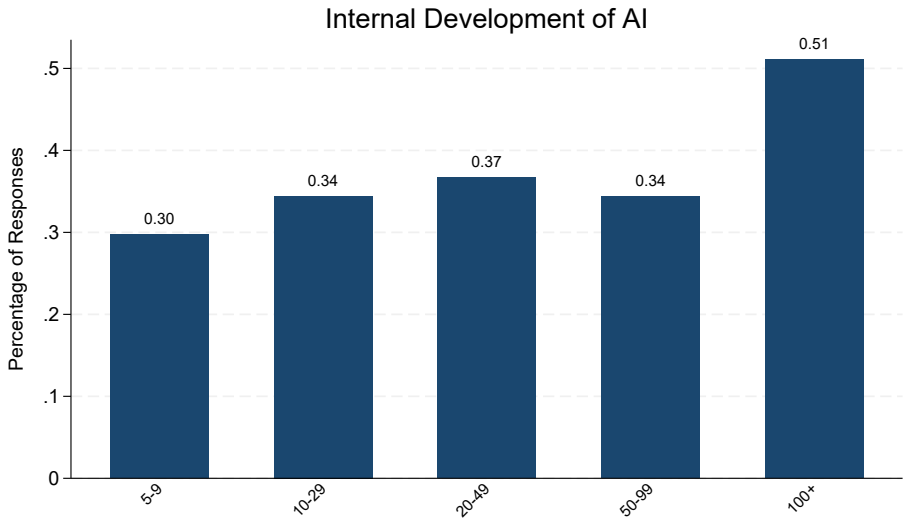

Note: Responses indicate whether the adopted AI system was developed internally. The sample is restricted to firms that report adopting at least one AI technology. The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Figure E.1.4: Distribution of Digitalisation Index

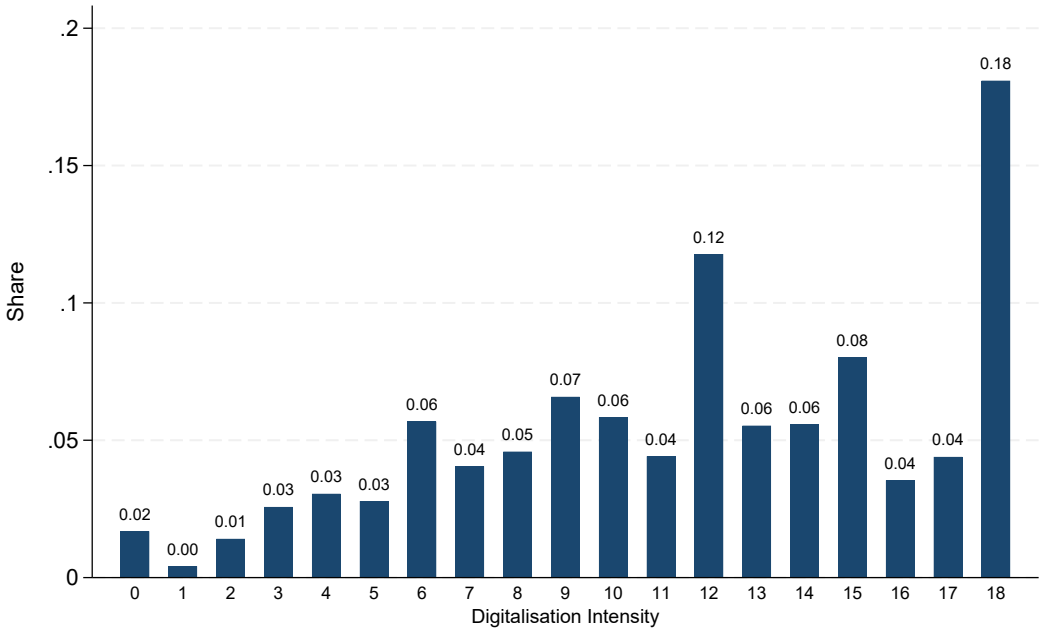

Note: The Digitalisation Index ranges from 0 to 18 and sums the degree of digitalisation (0–3) across six business functions: HR, Finance, Customer Feedback, Marketing, Supply Chain, and Production.

Figure E.1.5: Share of Employees with AI-Related Skills by Firm Size

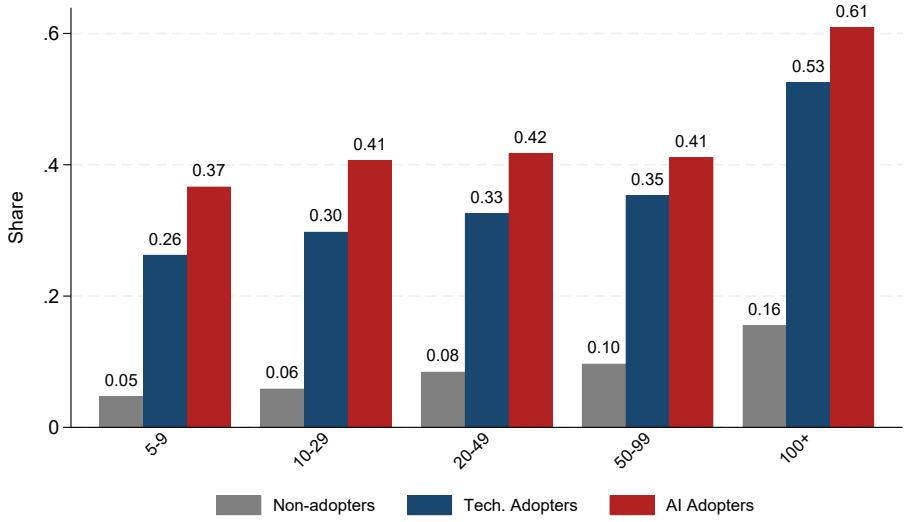

Note: The figure reports the weighted share of firms indicating that they employ workers with AI-related skills (e.g. programming, software development, and/or big-data management) by firm size and technology adoption status. These skills are relevant for the development, implementation, or maintenance of AI tools and systems. The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Other advanced technologies are AR: Augmented Reality, AGV: Automated Guided Vehicles, ASS: Automated Storage Systems, RFID: Radio Frequency Identification, Robotics, TS: Touch Screens/Kiosks, ADMS: Automated Decision-Making Systems. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

Figure E.1.6: Concerns about AI Adoption

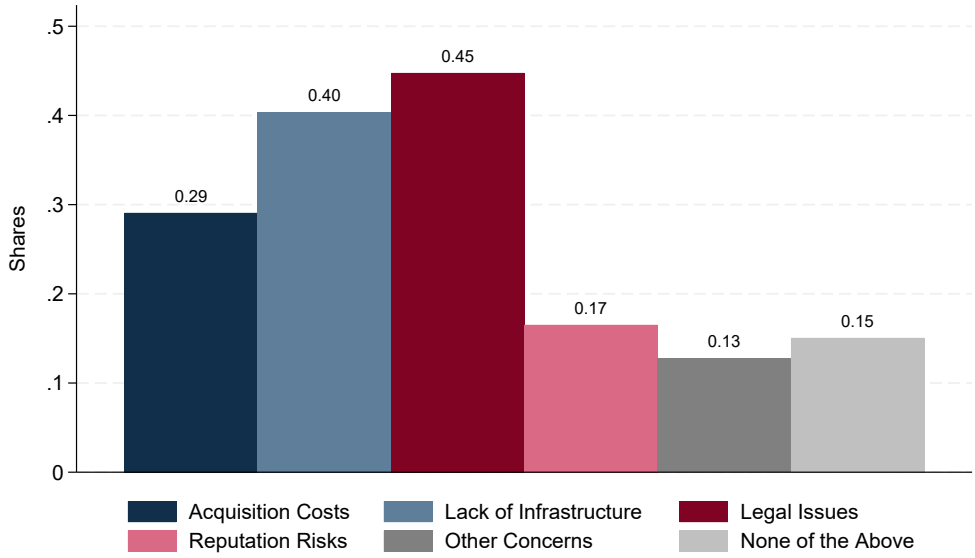

Note: The figure reports the share of responses to the question: "Please indicate whether the company has any of the following concerns about adopting AI." Multiple options could be selected. Statistics are weighted by the inverse probability of responding to the survey, as estimated in Appendix C.

E.2 Individual-Level Survey: Additional Figures

Figure E.2.7: Self-reported Adoption of Technologies in Workplaces

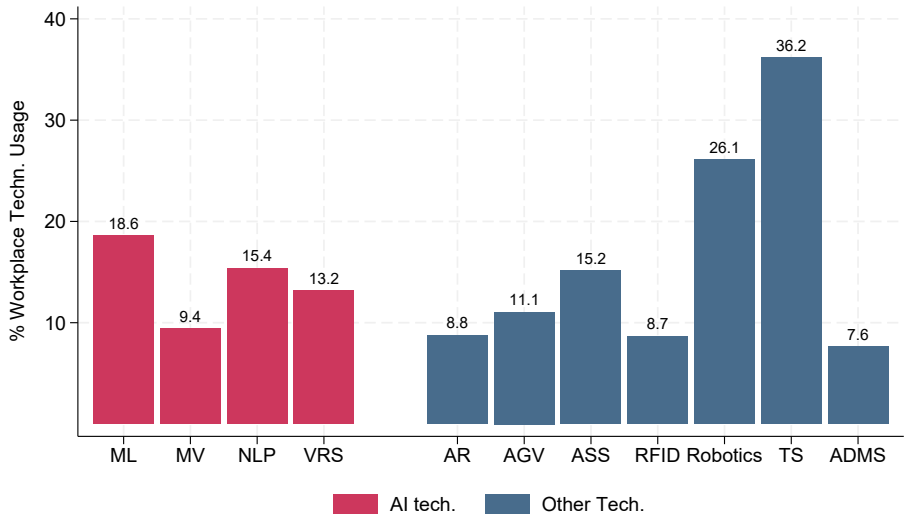

Note: The figure shows the share of workers reporting adoption of different technologies at the workplace. The AI technologies are ML: Machine Learning, MV: Machine Vision, NLP: Natural Language Processing, VRS: Voice Recognition Software. Other advanced technologies are AR: Augmented Reality, AGV: Automated Guided Vehicles, ASS: Automated Storage Systems, RFID: Radio Frequency Identification, Robotics, TS: Touch Screens/Kiosks, ADMS: Automated Decision-Making Systems. Statistics are weighted using survey weights.

Figure E.2.8: Self-Assessed AI Knowledge by Occupation

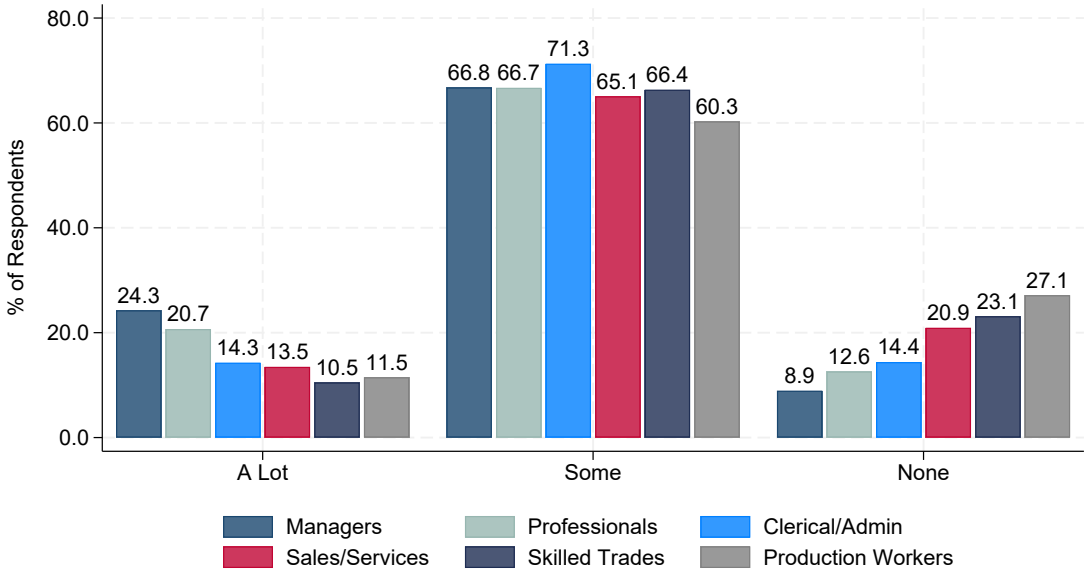

Note: Self-assessed AI knowledge based on the question: "How much would you say you personally know about artificial intelligence (AI)?" Responses are grouped by occupation. Statistics are weighted using survey weights.

F Main Results - Tables & Figures

Table F.1: Summary Statistics - TFP firm sample

	Main Estimation Sample		**TFP Sample**	
	Mean	Stand. Dev.	Mean	Stand. Dev.
Observations	2857		1773	
No of Employees	38.81	128.29	38.66	125.32
Revenue (Ml)	174.19	3381.83	130.33	738.75
Value Added (Ml)	38.75	266.94	37.77	211.25
Yearly Wage (th)	376.86	161.76	394.18	152.39
Hourly Wage	244.94	75.49	249.46	72.26
Firm Age	18.51	15.23	25.15	14.26
Exports (Ml)	72.36	2452.05	50.36	580.92
Imports (Ml)	107.71	5461.49	34.76	312.27
Physical Capital Stock (Ml)	8.70	171.49	9.91	103.18
Intangible Capital Stock (Ml)	0.03	0.61	0.04	0.79
Foreign Ownership	0.13	0.33	0.14	0.34
Digitalization	11.78	4.82	11.52	4.76
Cloud importance (0/1)	0.76	0.43	0.74	0.44
Union membership	0.58	0.22	0.62	0.20
Female Workers	0.32	0.26	0.29	0.24
Immigrant Workers	0.15	0.21	0.11	0.16
Age of Workers	41.37	7.88	43.40	7.37
Younger than 36	0.40	0.24	0.34	0.21
Between 36 and 54	0.38	0.18	0.39	0.16
Older than 55	0.22	0.19	0.27	0.19
No Secondary Education	0.19	0.18	0.20	0.17
Bachelors or Higher	0.33	0.30	0.27	0.26
Masters or Higher	0.12	0.19	0.09	0.15
STEM educated	0.26	0.26	0.31	0.26
University STEM educated	0.10	0.17	0.09	0.15
Managers	0.07	0.09	0.08	0.09
Professionals	0.31	0.33	0.27	0.31
Clerical/Administrative	0.12	0.18	0.11	0.17
Sales/Services	0.12	0.22	0.09	0.20
Skilled Trades	0.17	0.29	0.23	0.32
Production Workers	0.22	0.30	0.21	0.29

Notes: The table reports descriptive statistics for the main firm-level estimation sample and the subsample for which the TFP is estimated over the period 2008-2015. The TFP sample (n=1,773) is smaller than the main estimation sample (n=2,857) due to missing information on inputs required for TFP estimation or because the firm was not active before 2015. The comparison above shows that the composition of the sample for this robustness check is not radically different from the main sample.

Table F.2: Predictors of Technology Adoption (Extensive Margin) – *Panel A*

Dep. Variable:	Core AI					Other Tech.					Gen AI				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Firm characteristics															
Ln(employment)	0.054***	0.052***	0.054***	0.054***	0.049***	0.072***	0.070***	0.076***	0.070***	0.072***	0.075***	0.072***	0.070***	0.072***	0.112***
	(0.010)	(0.011)	(0.010)	(0.011)	(0.014)	(0.012)	(0.012)	(0.012)	(0.013)	(0.017)	(0.017)	(0.018)	(0.018)	(0.017)	(0.022)
Ln(value added per worker)	0.025	0.022	0.026	0.023	-0.005	-0.000	0.009	0.009	-0.006	0.011	0.031	0.029	0.033	0.022	0.016
	(0.021)	(0.022)	(0.023)	(0.025)	(0.026)	(0.021)	(0.022)	(0.021)	(0.023)	(0.030)	(0.028)	(0.027)	(0.026)	(0.028)	(0.046)
Ln(firm age)	-0.033**	-0.002	-0.006	0.000	0.012	-0.018	-0.018	-0.022	-0.023	0.029	-0.071***	-0.033	-0.046*	-0.057**	-0.009
	(0.016)	(0.017)	(0.017)	(0.017)	(0.028)	(0.018)	(0.019)	(0.019)	(0.019)	(0.033)	(0.022)	(0.024)	(0.024)	(0.024)	(0.038)
Foreign owned	0.028	0.010	0.012	0.005	-0.014	0.060*	0.057*	0.049	0.034	0.035	-0.072	-0.097**	-0.115**	-0.114***	-0.113**
	(0.029)	(0.029)	(0.029)	(0.029)	(0.033)	(0.034)	(0.034)	(0.034)	(0.034)	(0.037)	(0.047)	(0.045)	(0.047)	(0.042)	(0.049)
Exporting	0.043*	0.029	0.011	0.008	0.000	0.016	0.018	-0.002	-0.003	-0.002	0.076**	0.049	0.040	0.049	0.017
	(0.026)	(0.026)	(0.029)	(0.029)	(0.036)	(0.028)	(0.029)	(0.031)	(0.031)	(0.039)	(0.036)	(0.037)	(0.038)	(0.038)	(0.048)
Importing	0.022	0.003	-0.010	-0.015	-0.019	0.060*	0.067**	0.053*	0.053*	0.108**	0.115***	0.093**	0.062	0.064*	0.022
	(0.030)	(0.030)	(0.030)	(0.030)	(0.039)	(0.031)	(0.032)	(0.032)	(0.032)	(0.042)	(0.038)	(0.038)	(0.038)	(0.037)	(0.048)
Ln(capital intensity)	-0.006***	-0.001	-0.002	-0.002	-0.000	0.007***	0.006**	0.005*	0.005*	-0.000	-0.006*	-0.001	-0.001	-0.001	0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Ln(intangible per worker)	0.004	0.002	0.003	0.003	0.002	0.002	0.001	-0.000	0.000	0.002	-0.006	-0.010*	-0.010**	-0.009*	-0.001
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)
Employer Association (yes/no)	0.042*	-0.013	-0.021	-0.018	-0.011	-0.028	-0.030	-0.045	-0.040	-0.019	0.047	-0.002	-0.009	0.006	0.008
	(0.025)	(0.025)	(0.025)	(0.025)	(0.029)	(0.027)	(0.029)	(0.029)	(0.028)	(0.034)	(0.036)	(0.036)	(0.036)	(0.036)	(0.044)
Digitalisation index	0.104***	0.086***	0.081***	0.082***	0.085***	0.057***	0.060***	0.058***	0.059***	0.065***	0.115***	0.092***	0.087***	0.082***	0.077***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.014)	(0.013)	(0.013)	(0.013)	(0.013)	(0.015)	(0.016)	(0.016)	(0.016)	(0.015)	(0.019)
Cloud services	0.117***	0.084**	0.069**	0.064**	0.076**	0.029	0.035	0.040	0.043	0.044	0.167***	0.130***	0.121***	0.118***	0.098**
	(0.034)	(0.033)	(0.032)	(0.032)	(0.035)	(0.033)	(0.034)	(0.033)	(0.032)	(0.035)	(0.042)	(0.041)	(0.040)	(0.039)	(0.044)
Workforce characteristics		...	...	...	...		...	...	...	...		...	...	...	...
Sector FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Province FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Firm TFP FE					✓					✓					✓
Pseudo R^2	0.110	0.150	0.167	0.171	0.174	0.0617	0.0658	0.0840	0.0870	0.119	0.178	0.237	0.260	0.277	0.295
Observations	2857	2857	2857	2857	1773	2857	2857	2857	2857	1773	1440	1440	1440	1440	879

Notes: The dependent variable is an indicator equal to one if the firm reports adopting at least one of the technologies named in the columns and zero otherwise: Core AI (Machine Learning, Natural Language Processing, Machine Vision, or Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems). The unit of observation is a firm. Estimates are from logit model in equation 1 and reported as average marginal effects (AMEs). Workforce composition variables enter as shares; for each group of shares, the omitted category are workers aged 36–54 for age shares, male workers for female shares, natives for immigrant shares, workers with below-secondary education for education shares, and blue-collar workers for occupational shares. Exporting, importing, and foreign ownership are indicator variables. Firm TFP ventiles are fixed effects for ventiles of the firm-level total factor productivity distribution, estimated over the period 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015. The digitalisation index is standardised. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Robust standard errors are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.2 Cont.: Predictors of Technology Adoption (Extensive Margin) – *Panel B*

Dep. Variable:	Core AI					Other Tech.					Gen AI				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Firm characteristics	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...
Workforce composition															
Share of:															
No University, STEM		-0.121*	-0.154*	-0.176**	-0.227**		-0.187**	-0.092	-0.091	-0.134		-0.185*	-0.200*	-0.401***	-0.288**
		(0.071)	(0.079)	(0.086)	(0.096)		(0.076)	(0.083)	(0.091)	(0.100)		(0.096)	(0.108)	(0.116)	(0.128)
University, no STEM		0.223***	0.175**	0.173**	0.127		-0.097	0.001	-0.062	0.000		0.350***	0.314***	0.466***	0.605***
		(0.064)	(0.075)	(0.083)	(0.110)		(0.077)	(0.089)	(0.098)	(0.136)		(0.092)	(0.105)	(0.114)	(0.157)
University, STEM		0.248***	0.156**	0.099	0.177*		-0.030	0.081	0.044	-0.017		0.185*	0.164	0.183	0.001
		(0.066)	(0.074)	(0.080)	(0.095)		(0.084)	(0.092)	(0.099)	(0.117)		(0.095)	(0.111)	(0.123)	(0.151)
Age $<$ 39		0.116	0.163**	0.167*	0.156		0.006	0.024	0.092	0.078		0.046	0.094	0.088	0.329**
		(0.076)	(0.081)	(0.090)	(0.123)		(0.084)	(0.086)	(0.094)	(0.132)		(0.111)	(0.111)	(0.116)	(0.161)
Age $>$ 55		-0.056	0.015	0.039	0.093		0.020	0.007	0.020	0.113		-0.149	-0.128	-0.153	0.053
		(0.103)	(0.107)	(0.106)	(0.127)		(0.104)	(0.103)	(0.104)	(0.129)		(0.135)	(0.133)	(0.132)	(0.159)
Managers		-0.177	-0.223*	-0.240*	-0.318**		0.137	0.110	0.050	-0.029		0.315	0.175	0.155	0.788***
		(0.122)	(0.120)	(0.123)	(0.158)		(0.139)	(0.139)	(0.140)	(0.167)		(0.219)	(0.212)	(0.214)	(0.242)
White collar		0.044	0.033	0.072	0.139**		-0.051	-0.084	-0.106	-0.029		0.016	0.004	-0.080	0.030
		(0.048)	(0.054)	(0.057)	(0.068)		(0.057)	(0.060)	(0.065)	(0.080)		(0.066)	(0.073)	(0.074)	(0.092)
Immigrants				0.100	0.163				0.031	0.132				-0.353***	-0.516***
				(0.080)	(0.104)				(0.090)	(0.122)				(0.098)	(0.139)
Females				-0.117*	-0.002				0.103	-0.005				-0.182*	-0.045
				(0.066)	(0.082)				(0.074)	(0.095)				(0.097)	(0.121)
Union members				0.027	0.097				0.043	0.168				0.106	0.148
				(0.077)	(0.099)				(0.082)	(0.108)				(0.101)	(0.136)
Ln(hourly wage)				0.003	-0.014				0.155**	0.055				-0.009	-0.104
				(0.067)	(0.081)				(0.073)	(0.087)				(0.089)	(0.110)
Sector FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Province FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Firm TFP FE					✓					✓					✓
Pseudo R^2	0.110	0.150	0.167	0.171	0.174	0.0617	0.0658	0.0840	0.0870	0.119	0.178	0.237	0.260	0.277	0.295
Observations	2857	2857	2857	2857	1773	2857	2857	2857	2857	1773	1440	1440	1440	1440	879

Notes: The dependent variable is an indicator equal to one if the firm reports adopting at least one of the technologies named in the columns and zero otherwise: Core AI (Machine Learning, Natural Language Processing, Machine Vision, or Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems). The unit of observation is a firm. Estimates are from logit model in equation 1 and reported as average marginal effects (AMEs). Workforce composition variables enter as shares; for each group of shares, the omitted category are workers aged 36–54 for age shares, male workers for female shares, natives for immigrant shares, workers with below-secondary education for education shares, and blue-collar workers for occupational shares. Exporting, importing, and foreign ownership are indicator variables. Firm TFP ventiles are fixed effects for ventiles of the firm-level total factor productivity distribution, estimated over the period 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015. The digitalisation index is standardised. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Robust standard errors are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.3: Sector Heterogeneity in Adoption: Core AI vs. Other Advanced Technologies

	Core AI				**Other advanced technologies**			
	Manu.	Transp.	KIBS	Other	Manu.	Transp.	KIBS	Other
Firm characteristics								
Ln(employment)	0.055***	0.067***	0.080***	-0.006	0.117***	0.071***	0.099***	0.006
	(0.017)	(0.020)	(0.021)	(0.036)	(0.020)	(0.028)	(0.022)	(0.036)
Ln(value added per worker)	-0.022	0.050	0.003	0.013	0.065	-0.042	0.017	0.152*
	(0.037)	(0.038)	(0.046)	(0.067)	(0.054)	(0.045)	(0.030)	(0.080)
Ln(capital intensity)	-0.002	0.000	-0.007	0.001	-0.001	0.006	-0.001	0.002
	(0.003)	(0.004)	(0.007)	(0.009)	(0.004)	(0.004)	(0.007)	(0.010)
Ln(intangibles per worker)	0.001	0.002	0.013*	-0.017	-0.015*	0.000	-0.011	0.018
	(0.007)	(0.006)	(0.007)	(0.014)	(0.008)	(0.008)	(0.007)	(0.014)
Ln(firm age)	0.020	0.000	-0.020	0.009	-0.021	-0.017	-0.033	-0.019
	(0.024)	(0.034)	(0.033)	(0.047)	(0.033)	(0.038)	(0.035)	(0.047)
Foreign owned	0.029	0.019	-0.055	-0.073	0.079	0.082	-0.023	-0.086
	(0.042)	(0.042)	(0.064)	(0.111)	(0.069)	(0.055)	(0.065)	(0.115)
Exporting	-0.013	-0.034	0.027	0.040	-0.038	-0.081	0.087*	0.101
	(0.053)	(0.056)	(0.052)	(0.096)	(0.053)	(0.063)	(0.050)	(0.098)
Importing	-0.106*	0.199***	0.004	-0.093	0.052	0.176**	0.033	-0.141
	(0.055)	(0.070)	(0.052)	(0.082)	(0.051)	(0.078)	(0.055)	(0.086)
Employer association	-0.022	-0.023	-0.016	-0.179**	0.036	-0.099*	-0.071	-0.060
	(0.046)	(0.044)	(0.047)	(0.081)	(0.055)	(0.052)	(0.045)	(0.082)
Digitalisation index	0.111***	0.026	0.101***	0.086**	0.047**	0.037	0.058**	0.025
	(0.018)	(0.022)	(0.022)	(0.034)	(0.022)	(0.025)	(0.025)	(0.041)
Cloud services	0.048	0.074	0.127*	0.102	0.026	0.049	-0.029	0.136
	(0.043)	(0.048)	(0.070)	(0.087)	(0.048)	(0.058)	(0.066)	(0.094)
Workforce composition (shares)								
Age < 36	0.129	0.234	0.217	0.132	0.067	0.299	-0.054	0.168
	(0.151)	(0.170)	(0.137)	(0.195)	(0.166)	(0.199)	(0.139)	(0.255)
Age > 55	0.148	0.006	0.094	-0.069	0.182	0.242	-0.051	-0.266
	(0.168)	(0.203)	(0.198)	(0.270)	(0.190)	(0.221)	(0.186)	(0.321)
No University, STEM	-0.202**	-0.042	-0.155	-0.295	-0.224**	0.025	-0.196	-0.270
	(0.098)	(0.135)	(0.255)	(0.341)	(0.106)	(0.179)	(0.245)	(0.350)
University, no STEM	0.114	0.285**	0.276**	-0.252	-0.054	0.165	-0.166	-0.060
	(0.189)	(0.139)	(0.122)	(0.175)	(0.215)	(0.171)	(0.133)	(0.204)
University, STEM	0.450***	0.273	0.190*	-0.437	0.318	0.685***	-0.212	-0.255
	(0.158)	(0.169)	(0.112)	(0.463)	(0.200)	(0.223)	(0.132)	(0.583)
Managers	0.149	-0.263	-0.174	-0.070	0.228	0.246	0.108	0.990**
	(0.226)	(0.236)	(0.200)	(0.422)	(0.276)	(0.278)	(0.209)	(0.445)
White collar	-0.108	0.034	0.020	0.179	-0.207**	-0.189	-0.131	-0.051
	(0.090)	(0.107)	(0.163)	(0.111)	(0.103)	(0.137)	(0.128)	(0.121)
Province FE	✓	✓	✓	✓	✓	✓	✓	✓
Sector FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	963	741	721	257	963	741	721	257

Notes: The table reports average marginal effects (AMEs) from logit model in equation 1 estimated separately by broad sector. The specification in each column is the same as in columns 3 and 8 in Table F.2. The unit of observation is a firm. The dependent variable is an indicator equal to one if the firm reports adopting at least one of the technologies named in the columns and zero otherwise: Core AI (Machine Learning, Natural Language Processing, Machine Vision, or Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems). Regressions are weighted by the inverse probability of responding to the survey. All specifications include province fixed effects and detailed sector classification. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F.4: Predictors of Technology Adoption (Intensive Margin)— *Panel A*

Dep. Variable:	Number of Adopted Technologies														
Sample: Adopted Technology	Core AI					Other Tech.					Gen AI				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Firm characteristics															
Ln(employment)	0.352***	0.346***	0.343***	0.310***	0.489***	0.378***	0.350***	0.338***	0.329***	0.426***	0.460***	0.446***	0.451***	0.440***	0.666***
	(0.071)	(0.068)	(0.066)	(0.065)	(0.098)	(0.062)	(0.064)	(0.064)	(0.063)	(0.084)	(0.072)	(0.077)	(0.074)	(0.077)	(0.099)
Ln(value added per worker)	-0.202	-0.138	-0.065	-0.152	-0.072	-0.062	-0.052	0.010	-0.016	-0.012	-0.123	-0.131	-0.131	-0.146	0.053
	(0.144)	(0.147)	(0.148)	(0.131)	(0.170)	(0.149)	(0.142)	(0.148)	(0.138)	(0.171)	(0.168)	(0.164)	(0.159)	(0.146)	(0.180)
Ln(firm age)	-0.137	-0.093	-0.131	-0.125	-0.006	-0.205*	-0.064	-0.115	-0.075	0.053	-0.199	-0.159	-0.155	-0.136	0.187
	(0.145)	(0.137)	(0.132)	(0.133)	(0.202)	(0.121)	(0.114)	(0.116)	(0.116)	(0.174)	(0.161)	(0.177)	(0.170)	(0.174)	(0.185)
Foreign owned	0.729***	0.682***	0.612***	0.395**	0.185	0.541***	0.467***	0.431***	0.244	0.241	0.596***	0.518**	0.509**	0.397*	-0.119
	(0.233)	(0.195)	(0.192)	(0.190)	(0.217)	(0.193)	(0.168)	(0.165)	(0.155)	(0.186)	(0.211)	(0.204)	(0.199)	(0.204)	(0.220)
Exporting	0.210	0.206	0.180	0.127	-0.326	0.297	0.153	0.075	0.066	-0.359	-0.197	-0.175	-0.294	-0.268	-0.006
	(0.240)	(0.229)	(0.223)	(0.232)	(0.298)	(0.204)	(0.199)	(0.213)	(0.222)	(0.293)	(0.254)	(0.252)	(0.267)	(0.265)	(0.263)
Importing	-0.133	-0.140	-0.270	-0.242	0.121	-0.242	-0.180	-0.318	-0.377	-0.195	-0.242	-0.307	-0.404	-0.398	0.161
	(0.319)	(0.320)	(0.300)	(0.288)	(0.362)	(0.294)	(0.293)	(0.284)	(0.275)	(0.328)	(0.385)	(0.355)	(0.342)	(0.341)	(0.356)
Ln(capital intensity)	0.009	-0.005	-0.003	-0.002	-0.032	-0.028*	-0.016	-0.014	-0.016	-0.021	-0.022	-0.016	-0.013	-0.016	-0.089***
	(0.019)	(0.021)	(0.023)	(0.023)	(0.022)	(0.016)	(0.018)	(0.020)	(0.019)	(0.019)	(0.022)	(0.024)	(0.026)	(0.027)	(0.023)
Ln(intangible per worker)	0.051	0.050	0.041	0.044	0.023	0.069**	0.063**	0.052*	0.052*	0.017	0.033	0.020	0.003	0.008	-0.081**
	(0.038)	(0.036)	(0.032)	(0.031)	(0.038)	(0.033)	(0.031)	(0.031)	(0.029)	(0.035)	(0.052)	(0.049)	(0.044)	(0.043)	(0.040)
Employer association	0.073	0.127	0.053	0.117	-0.289	0.308	0.158	0.123	0.181	-0.182	-0.034	-0.118	-0.060	0.038	-0.579**
	(0.216)	(0.224)	(0.216)	(0.216)	(0.271)	(0.197)	(0.199)	(0.199)	(0.196)	(0.239)	(0.267)	(0.242)	(0.231)	(0.227)	(0.256)
Digitalisation index	0.209*	0.218*	0.306**	0.334***	0.248*	0.515***	0.450***	0.480***	0.473***	0.498***	0.327**	0.274**	0.281**	0.285**	0.285**
	(0.126)	(0.119)	(0.124)	(0.122)	(0.135)	(0.109)	(0.104)	(0.105)	(0.099)	(0.119)	(0.148)	(0.136)	(0.139)	(0.140)	(0.128)
Cloud services	0.030	0.125	0.166	0.131	-0.049	0.489*	0.425	0.412	0.323	0.158	-0.239	-0.220	-0.187	-0.250	-0.711**
	(0.288)	(0.297)	(0.275)	(0.276)	(0.290)	(0.285)	(0.291)	(0.281)	(0.271)	(0.271)	(0.388)	(0.400)	(0.353)	(0.351)	(0.283)
Workforce characteristics		...	...	...	...		...	...	...	...		...	...	...	...
Sector FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Province FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Firm TFP FE					✓					✓					✓
Pseudo R^2	0.0319	0.0451	0.0627	0.0708	0.591	0.0651	0.0811	0.0907	0.104	0.547	0.0443	0.0552	0.0780	0.0824	0.625
Observations	1104	1104	1104	1104	617	1343	1343	1343	1343	828	686	686	686	686	386

Notes: The table reports average marginal effects from a truncated Poisson model with mean parameter specified in equation (2). The unit of observation is a firm. The dependent variable is the number of distinct technologies adopted by the firm. The technologies are: Core AI (Machine Learning, Machine Vision, Natural Language Processing, Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems), and Gen AI. Columns (1)–(5) restrict the sample to firms that adopt at least one Core AI technology; columns (6)–(10) to firms that adopt at least one Other Technology; and columns (11)–(15) to firms that adopt at least one Gen AI technology. Workforce composition variables enter as shares. Exporting, importing, and foreign ownership are indicator variables. Firm TFP ventiles are fixed effects for ventiles of the firm-level total factor productivity distribution, estimated over the period 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015. The digitalisation index is standardised. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Robust standard errors are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.4 Cont.: Predictors of Technology Adoption (Intensive Margin) — *Panel B*

Dep. Variable:	Number of Adopted Technologies														
Sample: Adopted Technology	Core AI					Other Tech.					Gen AI				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Firm characteristics	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...
Workforce composition															
Share of workers:															
No University, STEM		-1.239*	-0.538	-0.325	0.337		-0.889	-0.794	-0.768	-0.065		-0.436	0.381	-0.493	-0.214
		(0.717)	(0.661)	(0.779)	(0.922)		(0.639)	(0.632)	(0.774)	(0.884)		(0.893)	(0.855)	(0.971)	(0.984)
University, no STEM		-0.829	-0.317	-1.183	-0.265		0.970	0.967	0.272	0.056		-1.131	-0.364	-0.220	-1.338*
		(0.669)	(0.744)	(0.743)	(0.818)		(0.628)	(0.728)	(0.688)	(0.787)		(0.903)	(0.841)	(0.873)	(0.730)
University, STEM		0.583	1.320*	0.538	0.270		1.616**	1.832**	0.940	0.956		0.603	1.358*	0.873	-1.472**
		(0.716)	(0.724)	(0.791)	(0.713)		(0.661)	(0.723)	(0.734)	(0.695)		(0.902)	(0.825)	(0.883)	(0.710)
Age $<$ 36		-0.568	-0.922	-0.065	-0.620		-0.387	-0.374	0.034	-1.387*		-0.751	-0.303	0.101	-0.300
		(0.811)	(0.757)	(0.798)	(0.920)		(0.818)	(0.731)	(0.710)	(0.804)		(0.784)	(0.779)	(0.884)	(1.056)
Age $>$ 55		-1.480	-1.770	-1.175	0.179		-1.903*	-1.637	-1.220	-1.578		-2.109*	-2.143	-2.005	-2.231*
		(1.064)	(1.094)	(1.026)	(1.020)		(0.993)	(1.015)	(0.904)	(0.963)		(1.209)	(1.322)	(1.290)	(1.148)
Managers		1.955*	2.185**	1.509	1.082		0.687	0.509	0.106	0.331		1.392	1.335	1.103	1.986*
		(1.026)	(1.018)	(1.019)	(1.322)		(0.893)	(0.909)	(0.876)	(1.224)		(1.342)	(1.315)	(1.299)	(1.186)
White collar		-1.097	-1.255**	-0.892*	-0.028		-0.661	-0.771	-0.070	0.501		0.139	0.125	0.012	1.277**
		(0.689)	(0.628)	(0.520)	(0.471)		(0.579)	(0.596)	(0.496)	(0.437)		(0.574)	(0.554)	(0.611)	(0.581)
Immigrants				1.569**	3.005***				2.020***	2.357***				0.198	1.840**
				(0.640)	(0.763)				(0.551)	(0.757)				(0.807)	(0.837)
Females				0.242	-0.312				-0.738	0.564				-0.759	0.176
				(0.573)	(0.694)				(0.553)	(0.597)				(0.672)	(0.761)
Unionization				0.636	-0.266				0.829	-0.770				1.255*	0.994
				(0.617)	(0.882)				(0.559)	(0.840)				(0.749)	(0.930)
Ln(hourly wage)				1.404***	1.558**				0.603	0.905*				0.293	1.264*
				(0.442)	(0.647)				(0.419)	(0.547)				(0.588)	(0.647)
Sector FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Province FE			✓	✓	✓			✓	✓	✓			✓	✓	✓
Firm TFP FE					✓					✓					✓
Pseudo R^2	0.0319	0.0451	0.0627	0.0708	0.591	0.0651	0.0811	0.0907	0.104	0.547	0.0443	0.0552	0.0780	0.0824	0.625
Observations	1104	1104	1104	1104	617	1343	1343	1343	1343	828	686	686	686	686	386

Notes: The table reports average marginal effects from a truncated Poisson model with mean parameter specified in equation (2). The unit of observation is a firm. The dependent variable is the number of distinct technologies adopted by the firm. The technologies are: Core AI (Machine Learning, Machine Vision, Natural Language Processing, Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems), and Gen AI. Columns (1)–(5) restrict the sample to firms that adopt at least one Core AI technology; columns (6)–(10) to firms that adopt at least one Other Technology; and columns (11)–(15) to firms that adopt at least one Gen AI technology. Workforce composition variables enter as shares. Exporting, importing, and foreign ownership are indicator variables. Firm TFP ventiles are fixed effects for ventiles of the firm-level total factor productivity distribution, estimated over the period 2008–2015 following Levinsohn and Petrin 2003 and Ackerberg et al. 2015. The digitalisation index is standardised. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Robust standard errors are reported in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.5: Deployment of AI Technologies by Business Function

	Business functions					
Technology	**CRM**	**SCM**	**FIN**	**HR**	**MKTNG**	**PROD**
Core AI						
Machine learning	0.097***	0.038***	0.150***	0.033	0.103***	0.185***
	(0.021)	(0.019)	(0.033)	(0.024)	(0.032)	(0.031)
Natural language processing	0.125***	0.050***	0.093***	0.097***	0.268***	0.145***
	(0.021)	(0.017)	(0.032)	(0.026)	(0.032)	(0.028)
Machine vision	-0.016	0.006	-0.040	0.015	-0.034	0.040
	(0.022)	(0.021)	(0.035)	(0.034)	(0.035)	(0.034)
Voice recognition	0.043	-0.026	-0.032	-0.019	0.030	0.000
	(0.028)	(0.018)	(0.034)	(0.030)	(0.038)	(0.031)
Other advanced technologies						
Augmented reality	0.019	0.016	-0.029	0.026	0.098***	0.027
	(0.030)	(0.026)	(0.035)	(0.036)	(0.045)	(0.045)
Automated guided vehicles (AGV)	-0.085***	-0.039***	-0.075***	-0.065***	-0.135***	-0.055
	(0.019)	(0.018)	(0.038)	(0.029)	(0.036)	(0.035)
Automated storage systems (ASS)	0.020	0.111***	0.001	-0.010	0.012	-0.034
	(0.023)	(0.034)	(0.039)	(0.029)	(0.037)	(0.030)
Radio-frequency identification (RFID)	0.035	0.018	0.032	0.021	0.003	0.000
	(0.026)	(0.022)	(0.033)	(0.029)	(0.033)	(0.029)
Robotics	0.008	0.006	0.075***	0.012	0.003	0.009
	(0.020)	(0.016)	(0.030)	(0.024)	(0.027)	(0.024)
Touch screens / kiosks	0.050***	0.052***	0.035	0.023	0.035	0.023
	(0.021)	(0.020)	(0.029)	(0.025)	(0.027)	(0.025)
Automated decision-making systems	0.068***	0.041	0.158***	0.170***	0.062	-0.012
	(0.036)	(0.030)	(0.057)	(0.054)	(0.049)	(0.039)

Notes: Entries are average marginal effects from specification (3). The dependent variable is an indicator for AI use in the business function listed in the column, regressed on indicators for adoption of the technology listed in the row. Acronyms: CRM = Customer Relationship Management; SCM = Supply Chain Management; FIN = Finance; HR = Human Resources; MKTNG = Marketing; PROD = Production. The sample comprises adopters of at least one technology ($N = 2{,}369$ firms; 12,828 observations). All specifications control for firm size and include sector and province fixed effects. Regressions are weighted by the inverse probability of responding to the survey, as estimated in Appendix C. Standard errors clustered at the firm level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F.6: Deployment of AI Technologies by Business Function, Including GenAI

	Business function (average marginal effects)					
Technology	**CRM**	**SCM**	**FIN**	**HR**	**MKTNG**	**PROD**
Core AI						
Machine learning	0.108***	0.083***	0.186***	0.072***	0.147***	0.144***
	(0.019)	(0.017)	(0.027)	(0.024)	(0.028)	(0.024)
Natural language processing	0.122***	0.064***	0.116***	0.108***	0.250***	0.135***
	(0.018)	(0.016)	(0.028)	(0.021)	(0.024)	(0.022)
Machine vision	0.061***	0.051***	0.069**	0.062**	0.093***	0.119***
	(0.023)	(0.019)	(0.034)	(0.027)	(0.034)	(0.025)
Voice recognition	0.089***	0.029	0.040	0.030	0.091***	0.050*
	(0.022)	(0.018)	(0.036)	(0.026)	(0.035)	(0.029)
Other advanced technologies						
Augmented reality	0.045*	0.046**	0.012	0.052*	0.109***	0.064**
	(0.024)	(0.020)	(0.038)	(0.030)	(0.040)	(0.032)
Automated guided vehicles (AGV)	-0.015	0.020	-0.009	-0.011	-0.053	0.018
	(0.023)	(0.018)	(0.040)	(0.031)	(0.043)	(0.032)
Automated storage systems (ASS)	0.026	0.071***	0.023	-0.004	-0.016	0.040
	(0.023)	(0.017)	(0.035)	(0.025)	(0.035)	(0.030)
Radio-frequency identification (RFID)	0.060***	0.036*	0.053	0.021	0.027	0.046
	(0.023)	(0.020)	(0.035)	(0.029)	(0.036)	(0.028)
Robotics	0.010	0.021	0.083***	0.015	-0.019	0.019
	(0.020)	(0.015)	(0.030)	(0.025)	(0.030)	(0.025)
Touch screens / kiosks	0.059***	0.038**	0.022	0.015	0.022	0.058**
	(0.020)	(0.017)	(0.029)	(0.024)	(0.029)	(0.024)
Automated decision-making systems	0.098***	0.081***	0.152***	0.128***	0.092*	0.024
	(0.027)	(0.023)	(0.046)	(0.032)	(0.050)	(0.039)
Generative AI						
Generative AI	0.265***	0.122***	0.234***	0.262***	0.665***	0.272***
	(0.061)	(0.043)	(0.053)	(0.056)	(0.070)	(0.049)

Notes: Entries are average marginal effects from separate logit regressions of an indicator for AI use in the business function listed in the column on an indicator for adopting the technology listed in the row. Acronyms: CRM = Customer Relationship Management; SCM = Supply Chain Management; FIN = Finance; HR = Human Resources; MKTNG = Marketing; PROD = Production. The sample comprises adopters of any technology ($N = 2{,}119$). Regressions involving Generative AI use only the 2024 wave where such question was included ($N = 1{,}084$). All regressions control for firm size and include sector and province fixed effects. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure F.1: Deployment of AI Technologies by Business Function, including GenAI

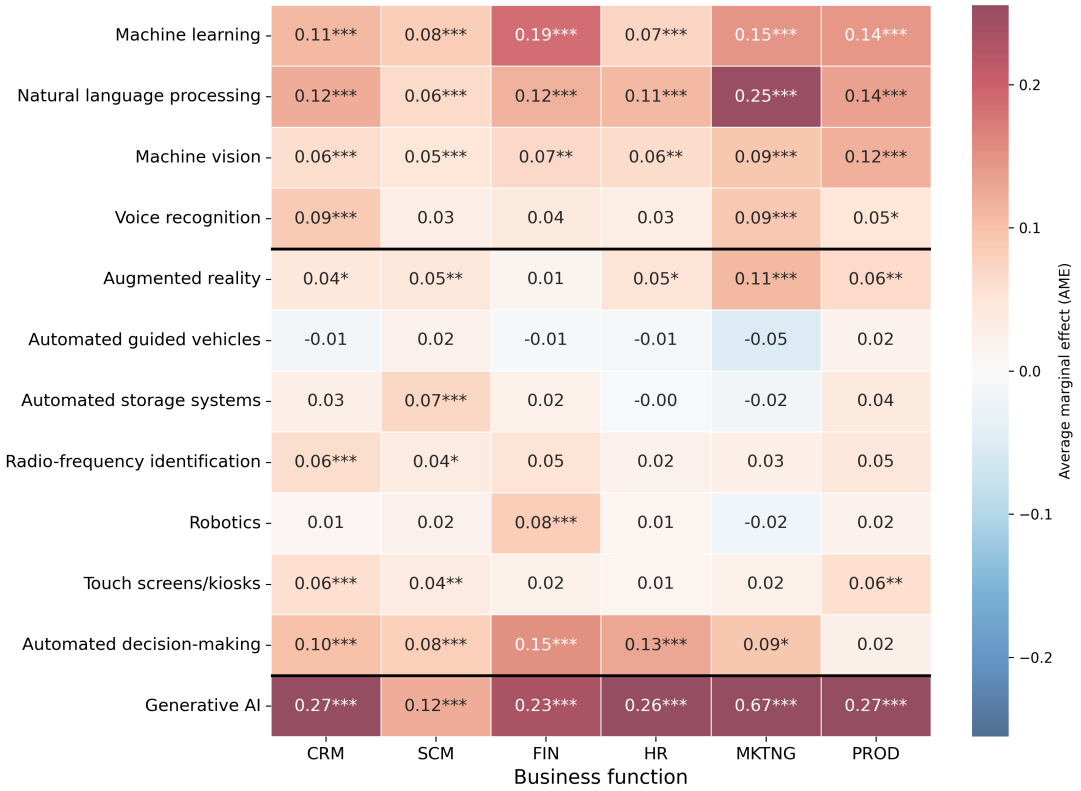

Notes: The figure is a visualization of Table F.6. Entries are average marginal effects from separate logit regressions of an indicator for AI use in the business function listed in the column on an indicator for adopting the technology listed in the row. Acronyms: CRM = Customer Relationship Management; SCM = Supply Chain Management; FIN = Finance; HR = Human Resources; MKTNG = Marketing; PROD = Production. The sample comprises adopters of any technology ($N = 2{,}119$). Regressions involving Generative AI use only the 2024 wave where such question was included ($N = 1{,}084$). All regressions control for firm size and include sector and province fixed effects. Robust standard errors in parentheses. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F.7: AI Adoption at the Workplace: Individual-Level Evidence

	Core AI		Other technologies	
	(1)	(2)	(3)	(4)
Worker characteristics				
Female	-0.126***	-0.092***	-0.067***	-0.038
	(0.023)	(0.032)	(0.026)	(0.036)
Age < 36	-0.003	0.026	0.001	0.034
	(0.026)	(0.038)	(0.029)	(0.041)
Age > 55	-0.049**	-0.061*	-0.070***	-0.056
	(0.024)	(0.035)	(0.026)	(0.038)
Immigrant	0.108***	0.081*	0.026	0.036
	(0.037)	(0.048)	(0.037)	(0.055)
Secondary, STEM	0.037	0.099	0.079	0.075
	(0.055)	(0.066)	(0.059)	(0.067)
Secondary, no STEM	0.000	0.013	0.086*	0.098
	(0.047)	(0.060)	(0.047)	(0.064)
Bachelor, STEM	0.047	0.080	0.143**	0.124
	(0.061)	(0.072)	(0.067)	(0.078)
Bachelor, no STEM	-0.024	0.102	0.094*	0.162**
	(0.050)	(0.068)	(0.051)	(0.076)
Master, STEM	0.175***	0.177**	0.168**	0.203**
	(0.061)	(0.082)	(0.069)	(0.090)
Master, no STEM	0.120**	0.099	0.099*	0.110
	(0.051)	(0.072)	(0.056)	(0.082)
Part-time	-0.101***	-0.090	-0.098***	-0.066
	(0.035)	(0.060)	(0.035)	(0.062)
Managers	0.264***	0.251***	0.141**	0.119*
	(0.055)	(0.071)	(0.060)	(0.071)
Professionals	0.215***	0.096*	0.055	0.059
	(0.036)	(0.049)	(0.044)	(0.055)
Clerical/Admin	0.193***	0.109*	0.069	0.045
	(0.044)	(0.061)	(0.052)	(0.065)
Sales/Services	0.045	0.006	0.024	0.089
	(0.040)	(0.069)	(0.048)	(0.071)
Skilled Trades	0.015	0.032	0.023	0.079
	(0.047)	(0.062)	(0.064)	(0.065)
Firm characteristics				
Ln(employment)		0.021**		0.043***
		(0.009)		(0.009)
Ln(value added per worker)		0.114***		0.068**
		(0.024)		(0.031)
Ln(firm age)		-0.005		0.003
		(0.020)		(0.021)
Foreign owned		0.032		0.026
		(0.036)		(0.038)
Exporting		-0.030		-0.010
		(0.052)		(0.050)
Importing		0.107*		0.106*
		(0.057)		(0.056)
Sector FE		✓		✓
Province FE		✓		✓
Pseudo R^2	0.090	0.168	0.023	0.119
Observations	2,250	1,000	2,250	1,000

Notes: The table reports average marginal effects from the Logit model specified in equation 4. The unit of observation is an individual. The dependent variable is an indicator equal to one if an individual reports technology use at the (most) current workplace, and zero otherwise. The technologies are: Core AI (Machine Learning, Machine Vision, Natural Language Processing, Voice Recognition Software), Other technologies (Augmented Reality, Automated Guided Vehicles, Automated Storage Systems, Radio Frequency Identification, Robotics, Touch Screens/Kiosks, Automated Decision-Making Systems). The sample is restricted to employed respondents whose workplace can be linked to firm-level data. Individual-level controls include gender, age group, education, STEM background, employment status, and occupation dummies. Firm-level characteristics refer to the worker's workplace. Effects for categorical variables are relative to omitted categories: males, workers aged 36–54, natives, individuals with below-secondary education, full-time employees, and production workers. Columns (2) and (4) exclude public sector employees, as value added is not available for public sector companies. Regressions are weighted using survey weights. Robust standard errors are reported in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table F.8: Individual Self-assessed AI Knowledge – Full Sample and Employed Subsamples

	(1)	(2)	(3)
Individual characteristics			
Female	-0.092***	-0.116***	-0.121***
	(0.016)	(0.019)	(0.029)
Age < 36	0.097***	0.107***	0.110***
	(0.016)	(0.019)	(0.027)
Age > 55	-0.069***	-0.076***	-0.111***
	(0.018)	(0.022)	(0.033)
Immigrant	0.000	0.034	0.077**
	(0.023)	(0.027)	(0.033)
Secondary, STEM	-0.058**	-0.018	0.000
	(0.025)	(0.034)	(0.038)
Secondary, no STEM	0.050**	0.057*	0.123***
	(0.020)	(0.029)	(0.036)
Bachelor, STEM	0.181***	0.172***	0.208***
	(0.043)	(0.060)	(0.071)
Bachelor, no STEM	0.060***	0.022	0.092*
	(0.022)	(0.032)	(0.050)
Master, STEM	0.277***	0.217***	0.177*
	(0.066)	(0.070)	(0.091)
Master, no STEM	0.147***	0.134***	0.104*
	(0.030)	(0.042)	(0.063)
Employment status (Omitted category: Full-time empl.)			
Part-time employment (status)	-0.046**		
	(0.023)		
Self-employed	0.017		
	(0.030)		
Unemployed	-0.017		
	(0.038)		
Student	0.128***		
	(0.021)		
Retired	-0.083**		
	(0.036)		
Homemaker	-0.029		
	(0.069)		
Other	-0.042		
	(0.030)		
Employed sample controls			
Part-time		-0.023	0.034
		(0.022)	(0.038)
Managers		0.199***	0.198***
		(0.041)	(0.040)
Professionals		0.135***	0.109**
		(0.035)	(0.043)
Clerical/Admin		0.163***	0.140***
		(0.037)	(0.044)
Sales/Services		0.060	0.064
		(0.039)	(0.048)
Skilled Trades		0.032	0.066
		(0.046)	(0.041)
Firm characteristics (employed sample)			
Ln(employment)			0.004
			(0.006)
Ln(value added per worker)			0.017
			(0.018)
Ln(firm age)			-0.016
			(0.014)
Foreign owned			0.070**
			(0.029)
Exporting			0.019
			(0.038)
Importing			0.045
			(0.037)
Sector and Province FE			✓
Pseudo R^2	0.119	0.129	0.251
Observations	3,357	2,250	1,000

Notes: The table reports average marginal effects from a Logit model where the dependent variable is an indicator equal to one if an individual reports high self-assessed knowledge of artificial intelligence, and zero otherwise. Column (1) uses the full sample and includes employment status controls. Columns (2)–(3) restrict the sample to employed respondents. Individual-level controls include gender, age group, education, STEM background, employment status, and occupation dummies. Firm-level characteristics refer to the worker's workplace. Effects for categorical variables are relative to omitted categories: males, workers aged 36–54, natives, individuals with below-secondary education, full-time employees, and production workers. Column (3) excludes public sector employees, as value added is not available for public sector companies. Regressions are weighted by survey weights. Robust standard errors are reported in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table F.9: AI Exposure and Robot Adoption

Exp. Measure	(1) Webb (2019)	(2) FRS (2018)	(3) FRS (2021)	(4) Eloundou et al. (2023)	(5) Engberg et al. (2024)
Dep. Var.: Robots adopted (survey)					
AI exposure	0.005	-0.050***	-0.033	-0.008	0.030
	(0.020)	(0.018)	(0.037)	(0.026)	(0.032)
Sector FE	✓	✓	✓	✓	✓
Province FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
R-squared	0.249	0.254	0.250	0.249	0.250
Within R^2	0.173	0.179	0.174	0.173	0.174
F-stat	15.30	16.36	15.33	15.27	15.38
F-KP	0.06	7.29	0.77	0.10	0.88
Observations	2,799	2,799	2,799	2,799	2,799

Notes: Each column reports the β from specification (5), estimated on separate models using a different AI exposure measure reported on columns (all standardized). Controls, measured in 2015, include firm size measured by employment quintiles, firm revenues measured in quintiles, firm age, indicators for foreign ownership, importing status, and exporting status, measures of exposure to software and robot automation based on Webb (2019), and the firm's baseline share of blue-collar workers. All regressions are weighted by the product of the estimated inverse probability of responding to the survey and firm size in 2015. Robust standard errors are reported in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table F.10: AI Exposure and Adoption of Other Technologies

Exp. Measure	(1) Webb (2019)	(2) Felten et al. (2018)	(3) Felten et al. (2021)	(4) Eloundou et al. (2023)	(5) Engberg et al. (2024)
Dep. Var.: Other tech. adopted (survey)					
AI exposure	0.024	-0.005	0.023	-0.005	0.041
	(0.018)	(0.017)	(0.033)	(0.025)	(0.029)
Sector FE	✓	✓	✓	✓	✓
Province FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
R-squared	0.192	0.191	0.192	0.191	0.193
Within R^2	0.141	0.140	0.140	0.140	0.141
F-stat	16.90	16.75	16.57	16.55	16.88
F-KP	1.73	0.09	0.46	0.03	2.07
Observations	2,799	2,799	2,799	2,799	2,799

Notes: Other advanced technologies are: Automated Guided Vehicles (AGV), Automated Storage Systems (ASS), Radio Frequency Identification (RFID), Touch Screens/Kiosks (TS), Robotics, Augmented Reality (AR), and Automated Decision-Making Systems (ADMS). Each column reports the β from specification (5), estimated on separate models using a different AI exposure measure reported on columns (all standardised). Controls, measured in 2015, include firm size measured by employment quintiles, firm revenues measured in quintiles, firm age, indicators for foreign ownership, importing status, and exporting status, measures of exposure to software and robot automation based on Webb (2019), and the firm's baseline share of blue-collar workers. All regressions are weighted by the product of the estimated inverse probability of responding to the survey and firm size in 2015. Robust standard errors are reported in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.