



Research paper

Artificial intelligence adoption and workplace training[☆]Samuel Muehlemann^{ID}*

LMU Munich, Munich, Germany

Institute of Labor Economics (IZA), Bonn, Germany

ROA Maastricht, Maastricht, The Netherlands



ARTICLE INFO

Dataset link: <https://doi.org/10.7803/371.1119.1.2.10>

JEL classification:

J23

J24

M53

O33

Keywords:

Artificial intelligence

Technological change

Automation

Apprenticeship training

Human capital

ABSTRACT

As artificial intelligence (AI) reshapes business processes, firms must adapt their training strategies to cultivate a skilled workforce. Using German establishment-level panel data from 2019 to 2023, this study analyzes how firms adjust their training strategies following AI adoption. Staggered difference-in-differences analysis shows that sustained AI adoption is associated with a 14% increase in new apprenticeships among training firms (intensive margin), but is not linked to the training decision (extensive margin). AI adoption is also associated with a modest increase in continuing training, with resources shifting toward high-skilled employees. The results align with AI as an automation innovation that reduces demand for simple skills as well as an augmentation innovation that increases demand for more advanced skills. The German dual apprenticeship system appears critical for firms aiming to build a future-ready workforce in the age of AI.

1. Introduction

As artificial intelligence (AI) diffuses across industries, managers face a critical dilemma: whether to train their existing workforce or acquire necessary skills externally by hiring new talent. This choice is complicated by AI's dual role as both an automation innovation, replacing routine and increasingly non-routine tasks, and as an augmentation innovation that enhances human capabilities in complex tasks (Autor et al., 2024; Autor, 2024). How firms balance training versus hiring will significantly influence labor-market dynamics and policy priorities in the AI era.

Germany, with its robust institutional support for vocational training, provides an ideal context for studying firms' internal skill formation strategies. Its established dual apprenticeship system offers a cost-effective channel for building a skilled future workforce internally. Instruction in vocational schools is publicly funded, while firms bear the cost of workplace-based training and pay apprentice wages that are substantially lower than those of fully qualified employees. This long-term investment contrasts with other strategies firms may use when AI creates urgent skill shortages, such as relying on employer-sponsored continuing training or external hiring to fill immediate gaps. By lowering the relative cost of the foundational "train" option, Germany's apprenticeship system may therefore tilt managerial decisions toward this form of strategic, long-term skill formation.

[☆] This study uses data from the BIBB Establishment Panel on Training and Competence Development 2019 to 2023. [doi:10.7803/371.1123.1.2.10](https://doi.org/10.7803/371.1123.1.2.10). I thank the Editor Daniela Puzzello and two anonymous referees for helpful comments, and the BIBB Research Data Center (FDZ), particularly Beatrice van Berk, Anett Friedrich, Felix Lukowski, and Sandra Hirtz, for their invaluable assistance with the data.

* Correspondence to: LMU Munich, LMU Munich School of Management, Geschwister-Scholl-Platz 1, D-80539 Munich, Germany.

E-mail address: muehlemann@lmu.de.

<https://doi.org/10.1016/j.jebo.2025.107206>

Received 11 October 2024; Received in revised form 24 July 2025; Accepted 16 August 2025

Available online 1 September 2025

0167-2681/© 2025 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Despite rapid AI diffusion, empirical evidence on its impact on firms' internal training strategies remains scarce. This study fills that gap by providing novel evidence on how sustained AI adoption affects workplace training. Using a representative panel of German firms from 2019–2023, we examine three key outcomes: the decision to offer apprenticeships, the number of new apprentices hired, and the intensity and skill-focus of continuing training.

We use the BIBB Establishment Panel on Training and Competence Development, containing representative firm-level data on training and AI adoption. Among SMEs, AI adoption rose from 4% in 2019 to 6.1% in 2023, whereas large firms (> 250 employees) saw a more substantial increase from 10% to 24%. To address the staggered timing of AI adoption across firms, we employ a staggered difference-in-differences approach (Callaway and Sant'Anna, 2021) to estimate its association with workplace training outcomes. This approach accounts for potential biases due to heterogeneous treatment effects across cohorts. We find that sustained AI adoption is associated with a 14% increase in new apprenticeships among training firms, but do not find a statistically significant link with the decision to offer apprenticeships. Second, AI adoption is associated with a modest increase in continuing training. However, the most notable finding concerns the composition of trainees: firms gradually shift training investments away from employees performing simpler tasks toward those performing high-skilled tasks, a shift that becomes pronounced two to three years after adoption. However, we find no significant impact of AI adoption on overall hiring, separation rates, or hiring composition by skill level.

Our findings illustrate how firms adapt their training strategies in response to technological change, offering important insights for policymakers seeking to facilitate effective AI integration. The positive association between AI adoption and the number of new apprenticeship positions in training firms suggests that the German apprenticeship system remains an important framework for skills development in the age of AI.

The remainder of this paper is structured as follows: Section 2 reviews the relevant literature. Section 3 presents the conceptual framework. Section 4 discusses the data and descriptive statistics. Section 5 expounds on the empirical strategy. Section 6 presents the results. Section 7 concludes the paper.

2. Literature review

The rapid diffusion of AI raises the question of how organizations adjust apprenticeship and continuing training strategies. To frame our empirical analysis, we review the following three strands of literature: (i) AI's dual role as both an automation and augmentation innovation, (ii) training responses to technological adoption, and (iii) the German apprenticeship context.

2.1. AI as automation and augmentation innovation

Evolving workplace technologies have consistently transformed labor markets by changing the allocation of tasks to labor and capital (Acemoglu and Autor, 2011). Earlier, automation primarily displaced routine tasks, but AI affects non-routine tasks at all skill levels, even those requiring high expertise (Autor, 2024). Feng and Graetz (2020) provide evidence from the U.S. that automation disproportionately targets tasks requiring substantial learned skills, contributing to employment polarization based on initial training requirements. The expanding reach of automation reinforces the need to understand its effects on the workforce.

Recent empirical research explicitly separates automation and augmentation effects. Using decades of patent and occupation data, Autor et al. (2024) find that while technology-driven automation erodes labor demand, augmentation innovations boost it by creating new work tasks. Their findings suggest, critically, that automation's labor-displacing effects have intensified in recent decades. By contrast, recent field studies of generative AI provide strong empirical evidence of its potential to augment worker productivity. Brynjolfsson et al. (2025) show that a generative AI tool used in customer service increased worker productivity by 15%, with the largest gains accruing to less-experienced and lower-skilled workers, thereby compressing the productivity distribution. This finding underscores that despite real displacement risks, AI's capacity to augment human expertise – particularly for less-experienced workers – is a defining aspect of current technological change.

2.2. Training responses to technological adoption

Empirical evidence at the firm-level is mixed regarding whether technological innovations complement or substitute for firm-sponsored training. Several studies provide strong evidence of complementarity. Bartel et al. (2007), studying earlier technologies, show that introducing computerized numerically-controlled machinery raised skill requirements and spurred corresponding training investments. More recently, Gathmann et al. (2024) find that German firms investing in digital technologies during the COVID-19 pandemic not only reported additional training needs but also acted on them by increasing the provision of employer-sponsored training.

However, other recent evidence points toward substitution, where technology adoption leads to a decline in training. A large cross-country study by Brunello et al. (2023) reveals that firms adopting advanced digital technologies (ADT) tend to reduce per-employee investment in training, arguing that ADT and training act as substitutes in production. Similarly, Heß et al. (2023) find that German workers highly exposed to automation are significantly less likely to receive firm-sponsored training, indicating active disinvestment in at-risk roles.

Further research suggests that training responses also depend on the technology type. Caselli et al. (2024) show that, in Italian firms, investment in information digital technologies (IDT) is associated with more IT-specific training and apprenticeship contracts, while investment in operational digital technologies (ODT) promotes more task-specific, non-IT training. These mixed findings highlight the need to isolate the effects of AI adoption within a clearly defined institutional context, motivating our empirical approach using German panel data.

2.3. The German apprenticeship context

Germany's dual apprenticeship system, combining firm-based training with part-time vocational schooling, remains the primary entry route to the labor market for roughly half of each youth cohort (BIBB, 2024). Its governance structure enables periodic revisions of occupational profiles through tripartite negotiations among employers, unions, and the federal government, allowing the incorporation of new technological requirements with relatively short lags (Jansen et al., 2017).

Historical evidence underscores the economic value of institutional flexibility in adapting rapidly to technological change. During Britain's Industrial Revolution, regions where guilds exercised less rigid control over craft training adapted to new technologies more swiftly (Zeev et al., 2017). Germany's modern reform mechanism similarly enables rapid updating of apprenticeship curricula to accommodate new technologies.

The system's design renders it a viable channel for building new skills. Programs typically last two to 3.5 years, with firms bearing part of the training costs, which apprentices' productive contributions partly offset (Muehlemann and Pfeifer, 2023). For firms, investing in apprentices can be an effective strategy to develop a talent pipeline and circumvent potentially high external hiring costs for specialized skills (Blatter et al., 2016). Clarifying the economic incentives driving firms' apprenticeship investments, especially in response to technological changes like AI, will be central to our conceptual framework.

3. Conceptual framework

This section synthesizes task-based theories of technological change with recent evidence on artificial intelligence (AI) to explain how firms adjust their skill strategies when adopting AI. Motivated by mixed evidence and the German context, the framework unfolds in three steps. First, we define the core technology channels – automation and augmentation – and their primary effects on skills demand, with a focus on how augmentation creates new work. Second, we analyze how firms respond to these new technologies through a two-layer matching process, leading to a strategic choice between internal skill development (“make”) and external hiring (“buy”). Third, we discuss the dynamic nature of this workforce adaptation over time.

3.1. How AI adoption affects skills demand

AI adoption alters a firm's demand for skills through two primary channels. The first is automation, which reallocates existing tasks from labor to capital. Its net effect on skill demand depends on the balance between productivity gains and labor displacement (Autor et al., 2024; Acemoglu and Restrepo, 2019). When productivity improvements are modest, so-called “so-so technologies” whose productivity gains marginally outweigh their costs, automation may reduce overall employment (Acemoglu and Restrepo, 2019).

In contrast, the augmentation channel complements human capabilities rather than simply replacing them. It enhances worker productivity by supporting decision-making and combining data-driven insights with worker experience, thereby enabling a broader set of workers to perform higher-stakes tasks previously reserved for elite experts (Autor, 2024). Critically, augmentation also drives a reinstatement effect by creating entirely new tasks in which labor holds a comparative advantage. These new tasks require specialized human expertise acquired through study and practical experience, forming a key driver of new labor demand (Autor et al., 2024). AI can also serve as an “innovation channel”, amplifying augmentation and reinstatement. By using it to create new AI-enabled task bundles, innovating firms raise the returns on complementary human skills (Rammer et al., 2022), thereby acting as architects of new work rather than passive technology adopters.

3.2. Skills strategy: matching, making, and buying skills

When AI disrupts task allocation, firms face two matching challenges: first, task-to-occupation matching, where AI reshapes which tasks define each occupation (Lazear, 2009); second, occupation-to-worker matching, where firms must source talent for these redefined roles. This dual matching challenge leads firms to strategically choose between developing skills internally (“make”) or hiring externally (“buy”), reflecting the varied firm responses – complementarity versus substitution – documented in previous literature. Although firms can “buy” skills externally, demographic-driven labor shortages and wage pressures in Germany limit this option. As a result, firms increasingly view the “make” strategy – developing skills internally – as a critical alternative.

The make strategy itself comprises two distinct pathways based on the target workforce and the timeline. For the incumbent workforce, firms can use short-term continuing training to upskill or reskill employees, bridging manageable gaps between existing competencies and new task requirements. For the future workforce, firms can invest in apprenticeships. This long-term strategy is especially potent in response to AI-driven augmentation. Creating new, specialized tasks requires a formal pipeline for building these novel competencies from the ground up. Apprenticeships represent a promising mechanism to train young workers to collaborate with and leverage AI tools, facilitating the widespread development of mid-level expertise in emerging skill domains (Autor, 2024). Young apprentices are economically attractive for AI integration, as direct training in firm-specific AI applications reduces adjustment costs associated with retraining experienced workers. Rather than becoming obsolete, apprenticeship curricula can be readily adapted to become a primary channel for creating the highly skilled, tech-savvy workforce of the future.

This leads firms to expand their apprenticeship programs along the intensive margin for three interconnected reasons. First, AI-driven augmentation creates entirely new task categories and occupational profiles, requiring additional apprentices beyond those needed for traditional core operations. Second, productivity gains from AI-human collaboration increase returns to scale for training

investments, economically justifying larger apprenticeship cohorts as each augmented worker manages more complex tasks and thus contributes more value-added. Third, the 2–3.5 year apprenticeship timeline aligns with firms' AI implementation timelines, making immediate program expansion optimal to ensure skilled workers are available when AI integration matures. The cost-effectiveness of apprenticeships becomes even more compelling when firms need to scale their workforce to capture AI-driven productivity gains.

The German institutional context strongly moderates this choice, particularly favoring apprenticeships as a “make” option. Specifically, the dual system's structure – combining publicly co-funded schooling with firm-based training, alongside wage compression and employment protection – lowers the net cost of training and limits the risk of post-training poaching (Dustmann and Schönberg, 2009; Muehleman and Pfeifer, 2023; Mohrenweiser et al., 2019). This institutional framework incentivizes firms to invest in apprenticeships not only to secure their future skill supply but also to strategically reduce long-term hiring costs and labor turnover, particularly when AI-induced skills gaps become too significant or costly to bridge through incumbent training alone.

3.3. Dynamic adjustment

When adopting AI, managers must evaluate the long-term costs and benefits of developing skills internally versus hiring externally, particularly as the technology creates both displacement and augmentation effects. Forward-looking firms recognizing AI's augmentation potential expand apprenticeship programs immediately after adoption, strategically aligning their human capital investments with AI maturity to maximize returns from AI-enabled productivity.

The emergence of augmentation-driven tasks, typically specialized and better remunerated (Autor et al., 2024), increases incentives to invest in apprenticeships due to higher expected returns and lower risks of poaching. Firms' dynamic capabilities (Teece, 2007) – including sensing technological opportunities, seizing resources for training, and reconfiguring workflows – shape their effectiveness in navigating the AI transition. As institutional processes such as curriculum approval and wage bargaining introduce lags, as previously discussed in the context of the German apprenticeship system's institutional flexibility, the full employment effects of AI adoption surface only gradually, reinforcing the value of analyzing firm responses over a longer time horizon.

Taken together, the framework predicts that German firms confronted with augmentation-oriented AI will prioritize internal skill development, using continuing training for incumbents and expanding and reorienting apprenticeships for their future talent pipeline. Expansion along the intensive margin – hiring more apprentices rather than merely updating training content – reflects firms' recognition that AI serves as a productivity-enhancing innovation rather than solely a labor-substituting technology. This shift enables firms to scale operations profitably by investing in a larger, more productive, and specialized workforce. The decision between workforce reductions and investment in human capital depends critically on the nature of the adopted technology—whether it represents marginal automation with uncertain returns or transformative augmentation with clear long-term potential.

4. Data and descriptive statistics

4.1. Data

This study utilizes the BIBB Establishment Panel on Training and Competence Development, a nationally representative, longitudinal dataset of German firms, from 2019 to 2023, to examine the relationship between AI adoption and firm-level training outcomes. The data are collected by the Federal Institute for Vocational Education and Training (BIBB) (Friedrich et al., 2023). The panel provides detailed information on firms' qualification structures, training behavior, organizational characteristics, and technology use, including AI adoption. The dataset is representative of the population of German firms. The sample size varies across years, averaging approximately 3000 establishments. In the 2019 survey wave, firms were first asked whether they use AI in their production, work, and business processes.¹

The survey also provides detailed information on firm-level training outcomes, including the number of employees participating in continuing training by task complexity levels (simple, skilled, and high-skilled) and apprenticeship programs. We also utilize data on wages, new hires, and employee separations, categorized by task complexity and sales growth, as additional outcome variables.

4.2. Descriptive statistics

Fig. 1 shows the AI adoption rate of German firms during 2019–2023, highlighting that adoption rates differ strongly by firm size. Among SMEs, AI adoption grew from 4% in 2019 to 6.1% in 2023; the increase was more substantial, from 10% in 2019 to 24% in 2023, for large firms with more than 250 employees.

¹ Specifically, they were asked: “Welche digitalen Technologien werden in Ihrem Betrieb für Produktions-, Arbeits- und Geschäftsprozesse genutzt?” [Which digital technologies are used in your company for production, work, and business processes?], and had the option to indicate the use of “Digitale Technologien, die auf dem Einsatz künstlicher Intelligenz und Maschinellen Lernen basieren, z. B. Deep Learning, Mustererkennung” [Digital technologies based on the use of artificial intelligence and machine learning, e.g., deep learning, pattern recognition]. In 2020 and 2021, the survey question was refined to distinguish between: (1) “Einsatz künstlicher Intelligenz und Maschinellen Lernen für physische Arbeitsprozesse, z. B. Deep Learning und Mustererkennung in Produktion und Wartung, Gebäudemanagement oder Pflege” [Use of artificial intelligence and machine learning for physical work processes, e.g., deep learning and pattern recognition in production and maintenance, building management, or healthcare], and (2) “Einsatz künstlicher Intelligenz und Maschinellen Lernen für nicht-physische Arbeitsprozesse, z. B. Deep Learning und Mustererkennung in Marketing, Beschaffung oder Personalwesen” [Use of artificial intelligence and machine learning for non-physical work processes, e.g., deep learning and pattern recognition in marketing, procurement, or human resources]. Owing to the low prevalence of AI adoption, particularly in the early years, these two categories were merged into a single binary variable, to denote any AI adoption in production, work, and business processes over the entire sample period.




Fig. 1. Artificial intelligence (AI) adoption rates by firm size (2019–2023)

Notes: This figure shows the mean adoption rates of AI technologies by year and firm size. Data are weighted by establishment weights.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

Table 1

Artificial intelligence (AI) introduction and discontinuation by firm size (2020–2023).

Source: BIBB establishment panel on training and competence development 2019–2023 long version

	Year			
	2020	2021	2022	2023
<i>Panel A: AI Introduction</i>				
Large Firms (250+ employees)	0.048 (0.214) [660]	0.041 (0.197) [641]	0.146 (0.353) [557]	0.100 (0.300) [350]
SMEs (<250 employees)	0.016 (0.125) [2,215]	0.019 (0.137) [2,380]	0.032 (0.177) [2,123]	0.039 (0.193) [1,466]
<i>Panel B: AI Discontinuation</i>				
Large Firms (250+ employees)	0.050 (0.219) [660]	0.038 (0.192) [641]	0.026 (0.159) [557]	0.096 (0.295) [350]
SMEs (<250 employees)	0.020 (0.141) [2,215]	0.023 (0.151) [2,380]	0.018 (0.132) [2,123]	0.021 (0.145) [1,466]

Notes: This table reports the mean rates of AI introduction and AI discontinuation by year and firm size. AI introduction refers to firms that newly adopted AI technologies in the given year. AI discontinuation refers to firms that stopped using AI technologies in the given year. Large firms are defined as having 250 or more employees, while SMEs have fewer than 250. Standard deviations are in parentheses, and sample sizes in square brackets. Data are weighted by establishment weights.

Table A.3 (Online Appendix) highlights key differences in workforce composition between AI-adopting and non-adopting establishments. AI adopters employ more workers on average and show a distinct qualification structure that evolves over time. While non-adopters maintained a stable skill distribution, AI adopters shifted markedly toward high-skilled employment between 2022 and 2023 (from 38.2% to 52.6%), with a corresponding decline in skilled-task workers (from 50.0% to 37.6%).

Table 1 presents the patterns of AI technology adoption and discontinuation compared to the previous year across different firm sizes from 2020 to 2023. Consistent with other data sources, the share of firms that discontinue the use of AI in production, work, and business processes is not negligible (Rammer et al., 2024). Panel A shows that large firms consistently adopt AI at higher rates than SMEs across all years, but Panel B reveals that large firms also discontinue AI use at higher rates than SMEs. The pattern is striking in 2023, when 9.6% of large firms, compared to 2.1% of SMEs, stopped using AI technologies. The higher rates of adoption as well as discontinuation among large firms suggest they may be more willing to experiment with AI technologies, yet ready to abandon implementations that do not satisfy expectations.

Table 2 shows apprenticeship training activities by AI adoption status from 2019 to 2023. The data reveal divergent patterns between the two firm types. Non-AI adopters consistently decline in apprenticeship participation over the five-year period, with training provision falling from 23.0% to 18.4%. This trend aligns with broader demographic shifts in Germany, particularly the decreasing number of school leavers entering the vocational training market (BIBB, 2024). AI adopters initially showed stronger

Table 2

Apprenticeship training by artificial intelligence (AI) adoption (2019–2023).

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

	Year				
	2019	2020	2021	2022	2023
<i>Panel A: Non-AI Adopters</i>					
Apprenticeship training	0.230 (0.421)	0.209 (0.407)	0.194 (0.395)	0.185 (0.389)	0.184 (0.388)
Number of new apprenticeships	0.239 (1.848) [1,856]	0.228 (1.325) [2,635]	0.223 (1.452) [2,774]	0.206 (1.596) [2,442]	0.214 (1.206) [1,635]
<i>Panel B: AI Adopters</i>					
Apprenticeship training	0.307 (0.463)	0.288 (0.454)	0.272 (0.446)	0.344 (0.476)	0.119 (0.325)
Number of new apprenticeships	0.583 (2.509) [152]	0.893 (8.536) [240]	0.487 (5.305) [247]	0.856 (2.750) [238]	0.465 (6.468) [181]

Notes: Standard deviations are in parentheses, and sample sizes in square brackets. Data are weighted by establishment weights. Apprenticeship training indicates whether firms offer apprenticeship training. New apprenticeship represents the number of new apprentices hired by firms in a calendar year.

Table 3

Apprenticeship training and technological change.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

	All Firms	AI Adopters	Non-AI Adopters
Integration of new firm-specific content	0.180 (0.384) [1,551]	0.395 (0.490) [174]	0.169 (0.375) [1,377]
Offering new training occupation	0.496 (0.500) [1,551]	0.683 (0.467) [175]	0.486 (0.500) [1,376]

Notes: This table reports the proportion of establishments implementing changes in their apprenticeship training in 2020. “Integration of new firm-specific content” indicates whether firms have established processes to regularly assess if new content (such as training necessitated by emerging technologies) should be incorporated into their existing apprenticeship programs owing to changing skill requirements. “Offering new training occupation” indicates whether firms have established processes to determine if apprenticeships should be provided in occupations not previously offered by the firm. Standard deviations are in parentheses, and sample sizes in square brackets. Data are weighted by establishment weights.

engagement with the apprenticeship system, maintaining higher participation rates through 2022. However, they experienced a notable decline in 2023, with the share of firms offering apprenticeships dropping to 11.9%. Despite this recent decrease in participation rates, AI-adopting firms consistently employed more apprentices on average across all years.

Table 3 reveals how establishments adapt their apprenticeship training to changing skill requirements, highlighting substantial differences between AI adopters and non-adopters. AI-adopting firms demonstrate a more dynamic approach to training innovation across two dimensions. First, they are more than twice as likely to have established processes for integrating new firm-specific content into existing apprenticeships (39.5% versus 16.9%).² Second, 68% of AI-adopting firms have processes to evaluate the potential for expanding apprenticeships into new occupational fields, compared to 49% of non-adopters. These differences persist even when controlling for structural factors, such as firm size and collective bargaining agreements.³ This pattern suggests that firms implementing AI technologies may be more responsive to evolving skill requirements and more systematic in adapting their training programs to satisfy emerging needs.

Table 4 examines continuing training patterns by AI adoption status from 2019 to 2023.⁴ The data reveal two consistent patterns. First, AI-adopting firms demonstrate a higher continuing training intensity in terms of training provision probability and the share of

² Note that this information was included only in the 2020 wave of the survey.

³ In a simple cross-sectional regression with these controls, the results in Table A.2 (Online Appendix) show that AI adopters are 15.6 percentage points more likely than non-adopters to have established processes to include new firm-specific content, and 18.7 percentage points more likely to consider new training occupations.

⁴ The survey defines continuing training as employee participation in internal or external courses, seminars, or workshops supported by the company through time off or financial contributions, excluding apprentices: “Haben Beschäftigte Ihres Betriebes im Jahr 2019 an sonstigen Weiterbildungsmaßnahmen in Form von internen oder externen Kursen, Seminaren oder Lehrgängen teilgenommen, die von Ihrem Betrieb durch Freistellung oder Kostenübernahme ganz oder teilweise gefördert wurden? Auszubildende bitte nicht berücksichtigen”. [Did employees of your establishment participate in other continuing training measures in 2019 in the form of internal or external courses, seminars, or training sessions fully or partially supported by your establishment through time off or cost coverage? Please do not include apprentices.]

Table 4

Continuing training indicators by AI adoption (2019–2023).

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

	Year				
	2019	2020	2021	2022	2023
<i>Panel A: Non-AI Adopters</i>					
Continuing training	0.688 (0.464) [1,856]	0.650 (0.477) [2,635]	0.498 (0.500) [2,774]	0.515 (0.500) [2,442]	0.618 (0.486) [1,635]
Share employees in continuing training	0.403 (0.382) [1,856]	0.376 (0.376) [2,635]	0.279 (0.360) [2,774]	0.300 (0.368) [2,442]	0.339 (0.364) [1,635]
Share trainees: simple tasks	0.039 (0.149) [1,574]	0.050 (0.176) [2,200]	0.046 (0.170) [1,977]	0.063 (0.200) [1,813]	0.065 (0.212) [1,333]
Share trainees: skilled tasks	0.644 (0.373) [1,574]	0.596 (0.395) [2,200]	0.622 (0.388) [1,977]	0.577 (0.392) [1,813]	0.621 (0.398) [1,333]
Share trainees: high-skilled tasks	0.302 (0.358) [1,574]	0.320 (0.381) [2,200]	0.314 (0.374) [1,977]	0.311 (0.364) [1,813]	0.299 (0.371) [1,333]
<i>Panel B: AI Adopters</i>					
Continuing training	0.749 (0.435) [152]	0.805 (0.397) [240]	0.602 (0.490) [247]	0.845 (0.362) [238]	0.701 (0.459) [181]
Share employees in continuing training	0.411 (0.371) [152]	0.457 (0.367) [240]	0.395 (0.422) [247]	0.541 (0.373) [238]	0.507 (0.425) [181]
Share trainees: simple tasks	0.021 (0.129) [131]	0.106 (0.282) [223]	0.028 (0.096) [209]	0.023 (0.107) [210]	0.033 (0.142) [157]
Share trainees: skilled tasks	0.569 (0.364) [131]	0.571 (0.394) [223]	0.665 (0.336) [209]	0.544 (0.386) [210]	0.451 (0.436) [157]
Share trainees: high-skilled tasks	0.407 (0.361) [131]	0.292 (0.361) [223]	0.260 (0.299) [209]	0.418 (0.378) [210]	0.514 (0.447) [157]

Notes: Standard deviations are in parentheses, and sample sizes in square brackets. Data are weighted by establishment weights. Continuing training indicates whether firms offer any continuing training. Share of employees in continuing training refers to the relative number of trainees by the level of task complexity (simple, skilled, high-skilled).

employees participating in continuing training activities. Second, the composition of training differs markedly between firm types. While non-adopters maintained relatively stable training distributions across task complexity levels through this period, AI adopters showed a pronounced shift toward high-skilled training. By 2023, over half (51.4%) of all continuing training in AI-adopting firms targeted high-skilled tasks, up from 40.7% in 2019, with a corresponding decrease in skilled-task training.

Overall, both apprenticeship and continuing training activities decreased during the pandemic, reflected by lower values in the years 2020 and 2021. However, continuing training activities increased again in subsequent years, although they did not reach pre-pandemic levels. Table 5 shows the share of COVID-19-affected firms that implemented short-time work policies specifically owing to the pandemic. Overall, about one-third of German firms introduced such measures, independent of AI adoption status, suggesting that AI adopters and non-adopters were affected similarly by the pandemic.⁵

Table A.4 (Online Appendix) shows employee separation patterns by AI adoption status from 2019 to 2023. The overall separation rates are broadly comparable between AI adopters and non-adopters. Likewise, Table A.5 (Online Appendix) shows no stark differences in the overall hiring rates, but there are some notable differences in the skill composition of new hires between AI adopters and non-adopters. By 2023, high-skilled workers constituted 55.7% of new hires in AI-adopting firms compared to just 19.4% in non-adopting firms, which represents a substantial shift from 2019 when high-skilled hires comprised 23.5% of new employment in AI-adopting firms. Correspondingly, the proportion of simple-task and skilled-task hires declined substantially among AI adopters, with skilled-task hiring dropping from 56.1% in 2019 to 32.5% in 2023.

⁵ Table A.8 (Online Appendix) also shows the results of a fixed-effects panel regression, confirming that firms that implemented short-time work policies owing to the pandemic did not differ in probability of adopting AI.

Table 5

Short-time work (*Kurzarbeit*) owing to the COVID-19 pandemic by artificial intelligence (AI) adoption status.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

	Mean	Standard Deviation	N
All Firms	0.336	0.472	3,019
Non-AI Adopters	0.336	0.472	2,772
AI Adopters	0.337	0.474	247

Notes: This table reports the prevalence of short-time work (*Kurzarbeit*) based on the 2021 survey wave. Data are weighted by establishment weights. The variable short-time work is binary, with 1 indicating that the firm implemented short-time work arrangements specifically owing to the pandemic.

Table A.6 (Online Appendix) presents wage levels, collective bargaining coverage, and sales growth for AI-adopting and non-adopting establishments from 2019 to 2023. AI-adopting firms consistently paid higher average monthly wages than non-adopters. Regarding collective bargaining coverage, no clear pattern distinguishes AI adopters from non-adopters. AI-adopting firms consistently reported a higher probability of positive sales growth compared to non-adopters, with particularly notable gaps in 2019 (61.6% vs. 38.8%), 2020 (56.7% vs. 36.8%) and 2023 (59.3% vs. 36.0%). This aligns with (Czarnitzki et al., 2023) who show that AI adoption increases productivity in German firms. The only exception was in 2021, when both groups experienced similarly reduced growth rates, likely owing to the pandemic.

In summary, AI-adopting firms differ from non-adopters by employing a higher proportion of high-skilled workers, offering more apprenticeships and continuing training, paying higher wages across all skill levels, and being more likely to report positive annual sales growth.

5. Empirical strategy

Our empirical strategy is designed to account for the staggered nature of AI adoption across firms. As a baseline approach, we consider a conventional two-way fixed effects (TWFE) model to control for time-invariant firm heterogeneity and common time shocks. While TWFE regressions can offer a valuable starting point, they are known to produce biased estimates in staggered adoption designs, potentially yielding inaccurate average treatment effects owing to negative weighting issues (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Sun and Abraham, 2021). Therefore, our main analysis and preferred specification uses the robust staggered difference-in-differences (DID) estimator developed by Callaway and Sant'Anna (2021).

5.1. Two-way panel fixed effects regression

The TWFE approach effectively accounts for time-invariant unobserved heterogeneity across firms and common temporal shocks. The empirical specification is given by:

$$Y_{it} = \alpha + \beta AI_{it} + \lambda_i + \delta_t + \epsilon_{it}$$

where Y_{it} represents the firm-level training outcomes of interest for firm i in year t , and AI_{it} is a binary indicator for whether firm i adopted AI in its production, work, and business processes in period t , λ_i and δ_t represent firm fixed effects and year fixed effects, respectively. ϵ_{it} captures the error term.

A key limitation of the TWFE approach is its assumption of homogeneous treatment effects over time, which may not hold given the rapid evolution of AI technology. Moreover, the decision to adopt AI is endogenous, as it depends on the firm's AI adoption strategy. Specifically, unobserved time-varying factors, such as contemporaneous changes in firm strategy or unexpected market shifts, may simultaneously influence both AI adoption and investment in training, leading to biased estimates. Therefore, the causal effect of AI adoption on the outcome variables can only be identified to the extent that it depends solely on observable firm characteristics and time-invariant unobserved factors at the firm level (λ_i). However, if time-varying unobserved factors, ϵ_{it} , related to the firm's training behavior, are correlated with AI adoption, the coefficient β will be biased.

5.2. Staggered AI adoption and heterogeneous treatment effects

To explicitly account for heterogeneous treatment effects, we apply the estimator proposed by Callaway and Sant'Anna (2021). The approach addresses key methodological challenges in staggered adoption settings where firms adopt AI technology at different points in time.⁶ The method estimates group-time average treatment effects, denoted as $ATT(g, t)$, which represent the average

⁶ Estimations have been carried out using the `csdid` package that is available for Stata 18, using the doubly robust estimation procedure discussed in Callaway and Sant'Anna (2021).

treatment effect on the treated for firms first adopting AI in period g . Moreover, event study parameters are estimated for periods before and after AI adoption to explore the dynamics of the treatment effects.

A key identification assumption is the parallel trends assumption, which requires that in the absence of AI adoption, treated and comparison firms would have followed parallel paths in their training outcomes. A potential violation of the assumption of parallel trends arises from anticipation effects. For example, firms might anticipate the need for increased worker skills when adopting AI and increase training investments before adoption. A pre-adoption increase in training could lead to a downward bias in our estimates of the post-adoption effect, as the additional increase in training following AI adoption would be smaller than expected. Such a behavior is conceptually related to Ashenfelter's Dip, observed when evaluating the effects of training programs on earnings where the earnings of participants decline in the period immediately before they enter the program (Ashenfelter, 1978). In this study, however, a potential pre-adoption increase in training would have the opposite effect and lead to a downward bias. While the assumption of parallel trends is fundamentally untestable, the insignificant pre-treatment estimates in the event study plots provide supporting evidence for its validity. Our control group comprises firms that never adopted AI in 2019–2023.

Finally, following the requirements of staggered DID estimators, which assume non-reversal of treatment, the analysis is limited to firms that continuously utilize AI after their initial adoption. This sample selection implies that our results may be subject to selection bias, as firms that successfully implement AI may differ systematically from those that attempt but fail to maintain its use. Consequently, the interpretation of our results must be limited to the association between AI adoption and training outcomes within the context of sustained AI implementation and cannot be generalized to all firms adopting AI.⁷ The following section addresses further potential threats to the internal validity of the staggered DID approach.

5.3. Threats to identification

While the staggered DID approach mitigates several threats to causal identification, potential endogeneity concerns remain. The identifying assumption is that, conditional on covariates and group-specific time trends and in the absence of AI adoption, the potential outcomes for training would have evolved similarly across AI adopters and non-adopters. However, several factors could potentially violate this assumption: unobserved time-varying factors, correlated with both AI adoption and workplace training decisions, could confound the estimates. For example, a change in management, with a new CEO prioritizing both technological innovation and skills development, could introduce a spurious correlation in the absence of a causal effect of AI on training behavior.

Furthermore, unanticipated changes in market conditions, regulatory shifts, or technological advances could impact AI adoption and training investments simultaneously. To mitigate some of these concerns, we include firm-size-specific time trends, time trends by collective bargaining status, and industry-specific time trends as robustness checks. Another potential concern is that the pandemic drove AI adoption and workplace training strategies. However, using information on whether a firm implemented short-time work policies as a measure of pandemic exposure intensity, we find no significant association between the use of short-time work policies and AI adoption (Online Appendix, Table A.8), thereby mitigating concerns about pandemic-related omitted variable bias.

An alternative approach would be to employ instrumental variables (IV) regression, provided an available instrument generates plausibly exogenous variation in AI adoption and is thus unrelated to a firm's training behavior. However, finding a valid instrument in our context is challenging. Czarnitzki et al. (2023) use AI adoption by other firms within the same two-digit industry as an adoption instrument. However, their analysis is cross-sectional, as they do not have repeated observations of AI adoption at the firm level. Replicating this approach within the context of the BIBB establishment panel did not produce a significant first-stage association between average industry AI adoption and AI adoption in a fixed effects panel regression. A possible explanation for this non-significant result is the relatively small number of industries in the BIBB panel ($N=8$) and the limited within-firm variation in AI adoption over time. More fundamentally, the AI adoption instrument might also impact competition in an industry when applied in a dynamic setting, which then would affect training outcomes, as increased competition at the industry level was previously found to be associated with workplace training (Bassanini and Brunello, 2011). Therefore, we rely on the staggered DID approach, which provides the most credible identification strategy given our data and research context, but do not interpret the results as strictly causal.

In summary, while our empirical estimates offer robust insights into the association between AI adoption and training outcomes, they primarily reflect the experiences of firms that sustain AI use over time. Consequently, caution should be exercised when generalizing these results to all firms, as those discontinuing AI, or adopting it only temporarily, may face different challenges and yield distinct training outcomes. Given the available data sources, our findings should be interpreted as representing the best available evidence on the association between AI adoption and training.

6. Results

Our main analysis relies on the staggered DID estimator proposed by Callaway and Sant'Anna (2021), which is robust to heterogeneous treatment effects across adoption cohorts. We present the aggregated average treatment effects on the treated (ATT) as our primary findings, using event-study plots to visualize the dynamics of these associations over time.

⁷ Reassuringly, AI discontinuation is not associated with changes in workplace training practices (Online Appendix, Table A.9).




Fig. 2. Event study plots showing dynamic associations between AI adoption and apprenticeship training

Notes: These figures show the event study estimates for the association between AI adoption and apprenticeship training (top) and the log number of new apprenticeships (bottom). Both models include firm and time-fixed effects, and controls for firm size, collective bargaining agreement status, and industry-specific trends. Vertical bars represent 95% uniform confidence intervals. The reference line at $t = 0$ marks the timing of AI adoption. Pre-treatment estimates ($t-3$ to $t-1$) are in gray.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

6.1. Apprenticeship training and continuing training

Our primary finding concerns the intensive margin of apprenticeship training. As shown in Table 6, we find a positive and statistically significant association between sustained AI adoption and the number of new apprentices. The overall ATT from our preferred specification (Column 3) indicates that AI-adopting firms have, on average, a 14% higher number of new apprentices post-adoption. For the extensive margin – the decision to offer apprenticeships at all – we find no statistically significant association.

The event study plots in Fig. 2 show that the coefficients prior to treatment are close to zero and not statistically significant. This aligns with the requirements of the parallel trends assumption that AI adopters and non-adopters do not exhibit different trends with regard to apprenticeship training. For the training decision (top panel), AI adoption is associated with a decrease in the probability of offering apprenticeship training in periods following AI implementation. Although not statistically significant, this result is consistent with the automation channel of our conceptual framework, which posits that some firms will have a lower demand for mid-level skills when new technologies primarily automate, rather than augment, existing tasks.

In stark contrast to the above, the analysis for the intensive margin (Fig. 2, bottom panel) reveals a positive and statistically significant jump in the number of new apprenticeships of 13% immediately upon AI adoption at period t . This effect remains positive and substantial in the subsequent two periods, before the estimate becomes statistically insignificant at $t+3$. Thus, the results suggest that while AI adoption may eventually lead some firms to discontinue apprenticeship training altogether, the firms that maintain apprenticeship programs significantly expand them.

Fig. 3 presents the event study estimates of the association between AI adoption and firm-level continuing training activities. Again, pre-treatment coefficients are not statistically different from zero. While individual coefficients in the treatment period and subsequent periods are positive and range around 3–5 percentage points, they are not statistically significant at conventional levels. While the aggregate post-treatment effect is 4.2 percentage points and statistically significant (Table 6), its economic magnitude is relatively modest given that the average training probability of non-adopters is approximately 60% (see Table 4).

The intensive margin analysis (bottom panel) also reveals a statistically insignificant coefficient for AI adoption for the share of employees who participate in continuing training. Empirical evidence from both specifications suggests that AI adoption cannot be fully characterized as an automation innovation, but instead aligns with the notion that AI also augments labor and in turn necessitates the continued provision of skills in the workplace.




Fig. 3. Event study plots showing dynamic associations between AI adoption and continuing training

Notes: These figures show the event study estimates for the association between AI adoption and continuing training (top) and the share of employees in continuing training (bottom). Both models include firm and time-fixed effects, and controls for firm size, collective bargaining agreement status, and industry-specific trends. Vertical bars represent 95% uniform confidence intervals. The reference line at $t = 0$ marks the timing of AI adoption. Pre-treatment estimates (t-3 to t-1) are in gray.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

We hypothesize that AI adoption induces a compositional shift, directing training resources relatively more toward higher-skilled employees, as strongly evidenced in Fig. 4. While the share of trainees performing simple and skilled tasks trends negatively over time, the most notable result is the relative increase in training for those in high-skilled roles. This association becomes statistically significant two years after adoption and grows to a substantial 9.5 percentage point increase in the share of high-skilled trainees by the third year. These findings align with the notion that AI technologies augment high-skilled tasks while substituting, albeit to a small degree, for simple and mid-level task execution of incumbent workers as firms progressively integrate AI technologies into their operational frameworks.

6.2. Separations, hiring, and wages

Beyond training adjustments, our conceptual framework suggests that AI adoption could increase churn, potentially leading to higher separation rates among lower-skilled workers driven by automation and higher hiring rates for higher-skilled workers for augmentation needs. External hiring provides an alternative strategy for acquiring new skills, as our conceptual framework implies. Therefore, we examine the dynamic association between AI adoption and firm hiring behavior, looking at both the overall hiring rate and the skill composition of new hires.

The results show that AI adoption is not significantly associated with changes in overall separation or hiring rates (Online Appendix, Figures A.2 and A.3). Likewise, we find no statistically significant changes in the skill composition of separations or hires following AI adoption. Although point estimates suggest some patterns, such as declining shares of simple-task workers and increasing shares of skilled workers, these effects do not reach statistical significance in any post-treatment period. Figure A.4 (Online Appendix) presents our findings on wage dynamics. Average wages across all workers show a positive trajectory following AI adoption, but the event study estimates are not statistically significant. Disaggregated by skill group, wages for simple-task and skilled workers show no significant changes following AI adoption. For high-skilled workers, wages show a positive trend after adoption, but point estimates are again not statistically significant.

These results suggest that AI adoption is associated, if at all, with modest increases in average wages, with limited evidence for differential effects across skill groups. Employer–employee linked data would be valuable for more detailed analysis to address within-firm heterogeneity, but are unavailable.




Fig. 4. Event study plots showing dynamic associations between AI adoption and share of trainees by task complexity

Notes: These figures show the event study estimates for the association between AI adoption and the share of trainees in simple tasks (top), skilled tasks (middle), and high-skilled tasks (bottom). All models include firm and time-fixed effects, and controls for firm size, collective bargaining agreement status, and industry-specific trends. Vertical bars represent 95% uniform confidence intervals. The reference line at $t = 0$ marks the timing of AI adoption. Pre-treatment estimates (t-3 to t-1) are in gray.

Source: BIBB establishment panel on training and competence development 2019–2023 long version.

6.3. Sales growth

Finally, we explore the association between AI adoption and firm performance using the available binary indicator for positive sales growth. Our conceptual framework hypothesizes that complementary human capital adjustments enhance performance gains from AI adoption. Theoretically, we expect firms that successfully align their training and hiring strategies with AI demands to realize more significant benefits, such as positive sales growth. The ideal empirical test compares performance trajectories among AI adopters based on the extent of their complementary adjustments. However, such an analysis faces significant empirical limitations in our data. The number of AI adopters per year and the potential limited variation in adjustment strategies within this group restrict the statistical power needed for a reliable comparison focused solely on the mediating role of adjustments. Consequently, we acknowledge this limitation in directly testing this mechanism, and proceed by examining the broader association between AI adoption and performance.

Figure A.5 (Online Appendix) displays the event study estimates comparing AI adopters to never-adopters using the staggered DID approach. The results show no differences between adopters and non-adopters prior to treatment and statistically insignificant coefficients for AI adoption, although the point estimates are positive and the economic magnitude is substantial, as on average, only about a third of non-adopting firms report positive sales growth in any calendar year (Online Appendix, Table A.6).

Table 6
AI adoption and training outcomes (ATT).

Variable	(1)	(2)	(3)
Training Participation	−0.0065 (0.0161) [11,357]	−0.0087 (0.0161) [11,357]	−0.0097 (0.0164) [11,357]
Log Apprentices	0.1081*** (0.0384) [3,568]	0.1199*** (0.0402) [3,568]	0.1333*** (0.0433) [3,568]
Continuing Training	0.0378* (0.0199) [11,357]	0.0374* (0.0199) [11,357]	0.0409** (0.0203) [11,357]
Share of Workers in Continuing Training	0.0225 (0.0194) [11,357]	0.0211 (0.0194) [11,357]	0.0191 (0.0195) [11,357]
Share of Trainees - Simple Tasks	−0.0150 (0.0116) [8,365]	−0.0133 (0.0118) [8,365]	−0.0124 (0.0122) [8,365]
Share of Trainees - Skilled Tasks	0.0021 (0.0167) [8,365]	−0.0001 (0.0169) [8,365]	−0.0035 (0.0174) [8,365]
Share of Trainees - High-Skilled Tasks	0.0156 (0.0135) [8,365]	0.0161 (0.0136) [8,365]	0.0174 (0.0138) [8,365]
Positive Sales Growth	0.0385 (0.0309) [11,357]	0.0377 (0.0311) [11,357]	0.0380 (0.0315) [11,357]

Notes: The table reports average treatment effects on the treated (ATT) from the Callaway and Sant'Anna (2021) difference-in-differences estimator. Column (1) presents baseline results without additional controls. Column (2) adds time trends by firm size and collective bargaining status. Column (3) further includes time trends by industry. Standard errors in parentheses. Number of observations in brackets.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

6.4. Robustness checks

In the previous analysis, we included time trends specific to industry, collective bargaining status, and firm size, to account for potentially heterogeneous trajectories across firms with different characteristics. To verify the robustness of our findings, we now examine how our results differ when excluding such time trends. Table 6 presents these robustness checks for various outcomes. For training participation, we find consistently negative, small, and statistically insignificant coefficients across all specifications. Regarding the number of new apprentices among training firms, we find robust, positive, and statistically significant associations that strengthen as controls are added, with the ATT increasing from 0.108 to 0.133. Likewise, continuing training exhibits positive and statistically significant associations across all specifications. The share of workers in continuing training shows consistent statistically insignificant coefficients, despite positive point estimates. When analyzing training by skill level, we find negative but insignificant coefficients for simple tasks, negligible coefficients for skilled tasks, and modest positive but insignificant coefficients for high-skilled tasks, which remain similar in size when adding additional time trends. Finally, AI adoption yields statistically insignificant coefficients for sales growth, although the point estimates are positive across all specifications. These results confirm the robustness of our main findings: AI adoption shows significant positive associations with the number of new apprentices among training firms and the provision of continuing training, with these associations remaining stable or becoming somewhat stronger as we account for time trends by firm size, collective bargaining status, and industries.

Figure A.6 (Online Appendix) presents event study estimates for apprenticeship training outcomes without heterogeneous time trends for additional covariates. For the extensive margin (top panel), pre-treatment coefficients are statistically insignificant, supporting the parallel trends assumption. Consistent with previous results that include heterogeneous time trends, the association between AI adoption and the probability of training apprentices becomes increasingly negative over time, but the point estimates remain statistically insignificant.

The intensive margin (bottom panel) is also consistent with previous results, as AI adoption is associated with an immediate and statistically significant increase in the log number of apprenticeships. These are consistent with those that include heterogeneous

time trends, except for the decline in the third year following AI adoption. Likewise, Figure A.7 and Figure A.8 in the Online Appendix show that the results for continuing training remain qualitatively similar when excluding time trends by firm size, collective bargaining agreement status, and industries.

Our DID analysis reveals that the association between AI adoption and apprenticeship training outcomes is qualitatively similar across adoption cohorts (Online Appendix, Table A.10). The notably more substantial effect for 2023 adopters (0.2892) may reflect the emergence of more advanced AI technologies, potentially including large language models (LLMs), that seem associated even more strongly with the number of new apprenticeships. While we acknowledge the limitations of strictly causal interpretations in this context, the stability of our estimates across multiple cohorts provides robust evidence of the positive association between AI adoption and the number of newly hired apprentices in training firms. For continuing training measures, we observe moderate heterogeneity in treatment effects across adoption cohorts (Online Appendix, Table A.11), with a positive association for adopters in 2020 and 2022, but null effects for 2021 and 2023. When considering the share of employees in training, the point estimates are positive except for 2021. In summary, while AI adoption appears to consistently increase apprenticeship numbers in training firms across cohorts (with stronger effects for recent adopters), its effects on continuing training are more varied and close to zero in some years. We find no evidence of a negative association between AI adoption and continuing training. The results for apprenticeship training, in particular, strengthen our confidence that AI adoption operates through similar mechanisms over time, even while its magnitude on some outcomes for continuing training shows temporal variation.

Independent of the estimation method, apprenticeships are strongly and positively associated with AI adoption. While the staggered DID approach addresses treatment heterogeneity, it does so at the expense of relying on a selected sample of firms that maintain AI adoption throughout the observation period. Nevertheless, the results remain qualitatively similar, even when compared to the TWFE estimates based on the entire sample, and thus include firms that eventually discontinued AI adoption during the observation window.

6.5. Comparison with two-way fixed effects estimates

For baseline comparison, we also estimate a conventional two-way fixed effects (TWFE) model. Recent econometric literature notes that TWFE estimates can be biased in staggered adoption settings with heterogeneous treatment effects (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021), which is why the staggered DID results are our preferred specification.

The TWFE estimates (Online Appendix, Figures A.9-A.11) point in the same direction as our main findings, but suggest a smaller magnitude for the association between AI adoption and number of new apprentices (approximately 7%). This is consistent with TWFE's potential to mask the full extent of the association.

7. Conclusion

In Germany, AI adoption in production, work, and business processes increased substantially in recent years, particularly among large firms. This study analyzed how increased AI adoption was associated with changes in firms' training behavior. We find that sustained AI adoption is associated with an approximately 14% increase in the number of apprentices in training firms.

Our findings suggest firms strategically use apprenticeships to internally develop AI-complementary skills, taking advantage of the German apprenticeship system's flexibility to update training content quickly. Apprenticeships thus enable firms to proactively and cost-effectively adjust workforce skills in response to AI-driven technological change. First, firms adopting AI typically already have established processes that allow for rapid integration of AI-relevant content into apprenticeship curricula. Second, apprenticeships provide a cost-effective way to build these skills due to subsidized vocational schooling and lower apprentice wages relative to skilled workers. Third, firms may perceive apprentices as particularly adaptable to new technologies because apprentices are typically younger and more tech-savvy. Thus, rather than relying solely on external hiring, firms expand apprenticeships strategically, investing early in a workforce specifically trained for future AI-enhanced productivity. We find no significant evidence that AI adoption changes hiring or separation rates. This contrasts with recent evidence from other contexts, such as the Netherlands, where automation increases worker separations (Bessen et al., 2025). This divergence may underscore the importance of Germany's labor market institutions, which appear to guide firms toward a strategy of internal adjustment through training rather than external adjustment via the labor market.

Regarding continuing training, AI adoption has a statistically significant but economically modest positive effect. Notably, continuing training resources shift from employees performing simpler tasks toward those in high-skilled roles, with effects emerging clearly 2–3 years post-AI adoption. Finally, the lack of a significant relationship between AI adoption and sales growth indicates that not all AI investments yield immediate financial returns.

Our results align with the conceptual framework differentiating AI's automation and augmentation roles. We find that AI operates on dual margins. As an automation innovation, AI reduces demand for simple tasks; as an augmentation innovation, it increases demand for advanced skills, evident in the expansion of apprenticeships in training firms and continuing training targeted at high-skilled workers. From a policy perspective, robust apprenticeship programs appear essential to support firms and workers in effectively adapting to rapid AI-driven technological change. The shift of training toward higher-skilled employees raises concerns about skill polarization, indicating the need for targeted policy support for workers in lower- and middle-skilled roles. Future research should examine longer-term outcomes, differentiate training-related effects across distinct AI technologies – especially generative AI – and ideally leverage linked employer–employee data to allow more precise estimation of skill demands and productivity impacts.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Gemini Advanced 2.5 pro to refine the text and improve readability. After using this tool/service, the author reviewed and edited the text and takes full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Online Appendix

The Online Appendix for this article can be found at <https://doi.org/10.1016/j.jebo.2025.107206>.

Data availability

This paper uses data from the BIBB Establishment Panel on Training and Competence Development 2011 to 2023 long, doi: <https://doi.org/10.7803/371.1119.1.2.10>. Access only through Research Data Center of BIBB.

References

- Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings. In: *Handbook of Labor Economics*. vol. 4, Elsevier, pp. 1043–1171.
- Acemoglu, D., Restrepo, P., 2019. Automation and new tasks: How technology displaces and reinstates labor. *J. Econ. Perspect.* 33 (2), 3–30.
- Ashenfelter, O., 1978. Estimating the effect of training programs on earnings. *Rev. Econ. Stat.* 60 (1), 47–57.
- Autor, D., 2024. Applying AI to Rebuild Middle Class Jobs. Working Paper 32140, National Bureau of Economic Research, Cambridge, MA.
- Autor, D., Chin, C., Salomons, A., Seegmiller, B., 2024. New frontiers: The origins and content of new work, 1940–2018. *Q. J. Econ.* 139 (3), 1399–1465.
- Bartel, A., Ichniowski, C., Shaw, K., 2007. How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills. *Q. J. Econ.* 122 (4), 1721–1758.
- Bassanini, A., Brunello, G., 2011. Barriers to entry, deregulation and workplace training: A theoretical model with evidence from Europe. *Eur. Econ. Rev.* 55 (8), 1152–1176.
- Bessen, J., Goos, M., Salomons, A., Van den Berge, W., 2025. What happens to workers at firms that automate? *Rev. Econ. Stat.* 107 (1), 125–141.
- BIBB, 2024. Datenreport zum Berufsbildungsbericht 2024: Informationen und Analysen zur Entwicklung der beruflichen Bildung. Bundesinstitut für Berufsbildung (BIBB), Bonn.
- Blatter, M., Muehleemann, S., Schenker, S., Wolter, S.C., 2016. Hiring costs for skilled workers and the supply of firm-provided training. *Oxf. Econ. Pap.* 68 (1), 238–257.
- Brunello, G., Rückert, D., Weiss, C.T., Wruuck, P., 2023. Advanced Digital Technologies and Investment in Employee Training: Complements or Substitutes?. IZA DP 15936, Institute of Labor Economics (IZA), Bonn.
- Brynjolfsson, E., Li, D., Raymond, L., 2025. Generative AI at work. *Q. J. Econ.* 140 (2), 889–942.
- Callaway, B., Sant’Anna, P.H., 2021. Difference-in-differences with multiple time periods. *J. Econometrics* 225 (2), 200–230.
- Caselli, M., Fourrier-Nicolai, E., Fracasso, A., Scicchitano, S., 2024. Digital Technologies and Firms’ Employment and Training. CESifo Working Paper 1105, Center for Economic Studies and ifo Institute (CESifo), Munich.
- Czarnitzki, D., Fernández, G.P., Rammer, C., 2023. Artificial intelligence and firm-level productivity. *J. Econ. Behav. Organ.* 211, 188–205.
- Dustmann, C., Schönberg, U., 2009. Training and union wages. *Rev. Econ. Stat.* 91 (2), 363–376.
- Feng, A., Graetz, G., 2020. Training requirements, automation, and job polarisation. *Econ. J.* 130 (631), 2249–2271.
- Friedrich, A., Gerhards, C., Mohr, S., Troltsch, K., Weis, K., 2023. BIBB Training Panel – An Establishment Panel on Training and Competence Development 2011 to 2021 long, GWA 1.0. Federal Institute for Vocational Education and Training, Research Data Center at BIBB, Bonn.
- Gathmann, C., Kagerl, C., Pohlen, L., Roth, D., 2024. The pandemic push: Digital technologies and workforce adjustments. *Labour Econ.* 89, 102541.
- Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. *J. Econometrics* 225 (2), 254–277.
- Heß, P., Janssen, S., Leber, U., 2023. The effect of automation technology on workers’ training participation. *Econ. Educ. Rev.* 96, 102438.
- Jansen, A., de Grip, A., Kriechel, B., 2017. The effect of choice options in training curricula on the demand for and supply of apprentices. *Econ. Educ. Rev.* 57, 52–65.
- Lazear, E.P., 2009. Firm-specific human capital: A skill-weights approach. *J. Political Econ.* 117 (5), 914–940.
- Mohrenweiser, J., Zwick, T., Backes-Gellner, U., 2019. Poaching and firm-sponsored training. *Br. J. Ind. Relat.* 57 (1), 143–181.
- Muehleemann, S., Pfeifer, H., 2023. Evaluating apprenticeship training programs for firms. IZA World Labor.
- Rammer, C., Doherr, T., Kinne, J., Lang, D., 2024. KI-Einsatz in Unternehmen in Deutschland: Strategische Ausrichtung und internationale Position. Technical Report, Bundesministerium für Wirtschaft und Klimaschutz (BMWK).
- Rammer, C., Fernández, G.P., Czarnitzki, D., 2022. Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Res. Policy* 51 (7), 104555.
- Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J. Econometrics* 225 (2), 175–199.
- Teece, D.J., 2007. Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strat. Manag. J.* 28 (13), 1319–1350.
- Zeev, N.B., Mokyr, J., Van Der Beek, K., 2017. Flexible supply of apprenticeship in the British Industrial Revolution. *J. Econ. Hist.* 77 (1), 208–250.