

// NO.26-001 | 01/2026

DISCUSSION PAPER

// MELANIE ARNTZ, MYRIAM BAUM, EDUARD BRÜLL, RALF DORAU,
MATTHIAS HARTWIG, BRITTA MATTHES, SOPHIE-CHARLOTTE MEYER,
OLIVER SCHLENKER, ANITA TISCH, AND SASCHA WISCHNIEWSKI

Low Barriers, High Stakes: Formal and Informal Diffusion of AI in the Workplace

Low Barriers, High Stakes: Formal and Informal Diffusion of AI in the Workplace*

Melanie Arntz^{3,5} Myriam Baum² Eduard Brüll^{4,7} Ralf Dorau²
Matthias Hartwig¹ Britta Matthes³ Sophie-Charlotte Meyer¹
Oliver Schlenker^{†3,6} Anita Tisch¹ Sascha Wischniewski¹

¹BAuA, ²BIBB, ³IAB Nuremberg, ⁴ZEW Mannheim, ⁵University of Erlangen-Nuremberg, ⁶Ifo Institute & LMU Munich, ⁷University of Mannheim

January 8, 2026

Abstract

Artificial intelligence (AI) is diffusing rapidly in the workplace, yet aggregate productivity gains remain limited. This paper examines the *dual diffusion* of AI – through both formal, employer-led and informal, employee-initiated adoption – as potential explanation. Using a representative survey of nearly 10,000 employees in Germany, we document a high extensive but low intensive margin of usage: while 64 percent use AI tools, only 20 percent use them frequently. This diffusion is strongly skill-biased and depends less on establishment and regional characteristics. While formality is associated with more frequent usage, training, AI-based supervision, and higher perceived productivity gains, it does not broaden access. These patterns suggest that widespread informal usage can coexist with limited productivity effects when complementary investments and organizational integration lag behind.

Keywords: artificial intelligence; AI; technology diffusion; formal and informal adoption; training; algorithmic management; productivity; inequality

JEL: O33, O32, J24, J81, C83

*This project was financially supported by the Federal Ministry of Labour and Social Affairs in Germany. We thank David Autor, Ludger Woessmann, Simon Wiederhold, and further participants of the CESifo Workshop on Skills, Tasks and Technologies in the AI Era for helpful comments, feedback, or conversations. †Corresponding author: Oliver Schlenker, Ifo Institute, Poschingerstr. 5, 81679 Munich, Germany. Email: schlenker@ifo.de.

1 Introduction

Artificial intelligence has entered the workplace at unprecedented speed. The public release of ChatGPT in late 2022 brought artificial intelligence into the mainstream and marked the beginning of a new phase of technological change. Generative AI is now widely regarded as a general-purpose technology with potential applications across industries, occupations, and tasks (Brynjolfsson et al., 2017; Dahlke et al., 2024; Nicoletti et al., 2020). Within only a few years, AI tools have become part of the daily work routines of millions of employees. This rapid diffusion has intensified debates about how AI will affect productivity, inequality, and the demand for skills.

AI diffusion differs fundamentally from earlier technological waves. In contrast to previous general-purpose technologies, which required substantial firm-level investments in hardware, infrastructure, and dedicated training, generative AI can often be accessed directly by individual workers through low-cost or free web-based interfaces. This accessibility enables a new mode of diffusion in which employees adopt and experiment with AI independently of formal organizational initiatives. As a result, AI spreads through two parallel channels: a formal, employer-led channel and an informal, worker-led channel. Firm-level data suggest that between one in ten and one in four German establishments use AI (Gerhards and Baum, 2024; Licht and Wohlrabe, 2024; Falck et al., 2024). Yet our representative worker-level evidence shows that almost two thirds of employees already use AI-based tools at least occasionally, and about one in five do so regularly. Notably, only around one third of users report that their main AI tool was introduced by their employer, highlighting the prevalence of informal, self-initiated use (Arntz et al., 2025).

Despite its rapid spread, AI has not yet produced measurable gains in aggregate productivity. The literature increasingly points to a “modern productivity paradox,” where improvements in task-level efficiency do not translate into economy-wide growth (Brynjolfsson et al., 2017). Evidence from earlier waves of digitalization suggests a similar pattern: new technologies often diffuse faster than firms can adapt their resource allocations, skills, workflows, and organizational practices (Polák, 2017; Schweikl and Obermaier, 2020; Nicoletti et al., 2020). Controlled experiments show that generative AI can raise productivity in specific tasks (Noy and Zhang, 2023; Brynjolfsson et al., 2025), yet aggregate indicators remain flat. This gap points to lagging complementary investments and incomplete organizational integration (Brynjolfsson et al., 2017; Acemoglu et al., 2022). Understanding how AI spreads within and across workplaces is therefore essential for explaining why its productivity potential is not yet fully realized.

The coexistence of formal and informal adoption channels may help explain the productivity paradox. Informal diffusion is likely to spread AI widely but often at low intensity and without structured training, supervision, or integration into work processes. Formal adoption, by contrast, can embed AI more deeply and productively but remains concentrated in particular firms and sectors. This asymmetry implies that although AI use is now widespread, much of its productivity potential remains unrealized. From a diffusion-theory perspective, such dual-channel processes resemble “epidemic” patterns in which adoption propagates unevenly across networks and regions (Dahlke et al., 2024). Comparable mechanisms have been documented for other emerging technologies where informal intermediaries support early experimentation but lack the institutional support required for system-wide diffusion (Colovic et al., 2025). The uneven spread of generative AI is also reflected in individual-level adoption patterns. Cross-country evidence shows that younger, better-educated, and higher-income workers are more likely to use AI tools, particularly in knowledge-intensive occupations (Bick et al., 2024; Humlum and Vestergaard, 2024). Women remain less likely to use generative AI across nearly all sectors (Otis et al., 2024; Aldasoro et al., 2024; Chugunova et al., 2026), which is to a large extent explainable by differences in familiarity. Personality traits such as openness and agreeableness further mod-

erate individual uptake (Kaya et al., 2024). These patterns suggest that informal diffusion may reinforce existing inequalities in skills and opportunities, while formal adoption may amplify productivity advantages among firms already investing in complementary assets. Understanding the determinants and consequences of formal and informal AI adoption is therefore of high policy relevance.

This paper provides the first representative worker-level evidence on how AI diffuses through both formal and informal channels inside workplaces. To measure how AI enters firms, we designed the linked employer-employee *Digital Transformation and the Changing World of Work* (DiWaBe 2.0) survey. The 2024 survey covers nearly 10,000 employees across all sectors of the German economy and is linked to rich administrative records, offering a highly detailed view of how AI is adopted, used, and governed within and across establishments. This setup allows us to identify how individual, establishment, and regional characteristics shape AI adoption, and how the formality of implementation conditions outcomes such as usage intensity, training, AI-based supervision, and perceived productivity. By distinguishing employer-led (*formal*) from employee-led (*informal*) adoption, we are able to document a dual diffusion pattern that combines top-down coordination with bottom-up experimentation.

We find that AI diffusion is widespread: Roughly two thirds of employees report using AI-based tools. Yet only a minority (about 20 percent) use them frequently. Usage is also skill biased: Higher-skilled workers are substantially more likely to adopt and intensively use AI, whereas establishment and regional factors play a smaller role. This suggests that current diffusion reflects individual initiative and skill endowment more than organizational strategy. Formal, employer-led adoption does not expand the pool of users, but it does determine how deeply AI becomes embedded in daily work. Workers using employer-provided tools engage more intensively, participate in more training, and report stronger productivity gains, but they also face more AI-based supervision.

With these findings, we make three main contributions. First, by introducing the concept of *dual diffusion*, we are able to empirically distinguish between formal (employer-led) and informal (employee-led) adoption of AI at the workplace. While prior research has investigated firm-level diffusion mechanisms and the spatial or network dependencies of AI adoption (Dahlke et al., 2024), and others have analysed how innovation intermediaries facilitate the institutionalisation of digital transitions across organisations (Colovic et al., 2025), evidence on the *micro-foundations* of diffusion – how individual workers adopt, adapt, and integrate AI into their tasks – remains scarce. Recent surveys begin to document such micro-level patterns in specific professional populations (e.g., among researchers in Chugunova et al. (2026) and Van Noorden and Perkel (2023)), but evidence for a representative sample of employees across the economy remains limited. Our approach extends the literature by demonstrating that diffusion processes unfold both through organisational infrastructures but also through bottom-up experimentation by workers, thereby adding a new layer to the study of technological diffusion.

Second, we provide new representative evidence on the prevalence, intensity, and correlates of each diffusion mode. This worker-level evidence complements recent studies that examine firm-level AI adoption and its relationship to skill requirements, training, and employee changes (Muehlemann, 2025; Gerhards and Baum, 2024; Gualandri and Kuzior, 2024). By focusing on individual-level adoption channels, we show that the formality of AI implementation strongly conditions the intensity of usage and the extent of worker training, AI-based supervision, and self-reported productivity gains.

Third, we contribute to the broader discourse on innovation systems and the governance of emerging technologies (Lundvall, 1992; Fagerberg et al., 2005). Our findings suggest that innovation systems traditionally designed for firm-led technological change may miss informal, worker-led diffusion. By identifying this new hybrid mode of technological change, we offer an

explanation for the persistent productivity paradox of AI – informal diffusion spreads technology broadly but shallowly, whereas formal adoption embeds it deeply but selectively – and derive implications for innovation governance, skill policy, and organizational learning.

The remainder of the paper is structured as follows. Section 2 develops the conceptual framework of dual diffusion in the age of AI, situating it within diffusion theory and innovation-system perspectives. Section 3 describes the data and provides descriptive evidence on the prevalence and intensity of AI use, distinguishing between formal and informal adoption. Section 4 outlines the empirical strategy, followed by Section 5 which presents the main findings on usage patterns, training, supervision, and perceived productivity. The last section concludes with implications for policy and future research.

2 The Concept of Dual Diffusion of AI and Its Implications

Technological diffusion is a central mechanism through which innovation reshapes economies and workplaces. Classical diffusion theory views adoption as a process of communication and learning within social systems (Rogers, 2003). In this view, technological uptake depends not only on technical feasibility but also on the social networks and information channels that connect potential adopters. Building on evolutionary economics, subsequent contributions have modelled diffusion as a dynamic process of search, imitation, and capability accumulation (Nelson and Winter, 1982; Battisti and Stoneman, 2003, 2010). Across these perspectives, the locus of adoption has typically been the firm, which acts as the primary organisational interface between innovation and production. Such frameworks have provided powerful explanations for past waves of technological change, yet they are increasingly challenged by the rise of digital and data-driven technologies that diffuse beyond organisational boundaries.

Recent evidence highlights that digital technologies, and AI in particular, diffuse through networked and clustered patterns across firms and regions. Empirical studies document "epidemic" effects whereby adoption is influenced by spatial proximity, supply-chain connections, and professional networks (Dahlke et al., 2024). While prior research has investigated firm-level diffusion mechanisms and the spatial or network dependencies of AI adoption (Dahlke et al., 2024), and others have analysed how innovation intermediaries facilitate the institutionalisation of digital transitions across organisations (Colovic et al., 2025), evidence on the *micro-foundations* of diffusion – how individual workers adopt, adapt, and integrate AI into their tasks – remains scarce. At the same time, the institutional environment and intermediaries that translate technological opportunities into practice play a crucial role in shaping diffusion outcomes (Colovic et al., 2025). Firm-level analyses show that AI adoption is linked to organisational learning and training structures (Muehlemann, 2025). While these studies underscore the importance of system-level factors, they continue to conceptualise diffusion primarily as a firm-led process. In contrast, early German evidence from SOEP 2020 (Büchel and Monsef, 2024; Giering and Kirchner, 2021) and more recent workplace surveys (Kunze et al., 2025; Hall, 2024) suggests that the increasing accessibility of AI tools enables workers to adopt technologies directly, suggesting that an additional layer of diffusion operates at the individual level.

AI diffusion today operates through two interrelated channels – a formal and an informal one. We define *formal diffusion* as adoption that occurs through the top-down introduction of AI by employers, typically accompanied by training, supervision, or task redesign. By contrast, *informal diffusion* captures bottom-up, worker-initiated use of AI tools that are easily accessible without organisational investment or authorisation. This distinction aligns with research on user innovation, which emphasises that individuals frequently adapt or develop technologies outside formal organisational frameworks (von Hippel, 2006). It also resonates with theories of absorptive capacity, which view adoption as contingent on the skills and learning capabilities

of both individuals and firms (Cohen and Levinthal, 1990). In coordinated innovation systems (Lundvall, 1992), firm-led channels remain essential for embedding new technologies, but informal adoption allows experimentation to occur even in the absence of institutional coordination. Firm-level analyses emphasise that AI diffusion depends strongly on organisational learning environments and sectoral interlinkages (Falck et al., 2024; Na et al., 2023; Gerhards and Baum, 2024), while within firms, network and spill-over effects reinforce adoption among employees once AI has been introduced (Dahlke et al., 2024), highlighting how formal and informal mechanisms interact within establishments. Together, these mechanisms form a *dual diffusion regime* that is distinctive of the AI era: formal channels facilitate depth and integration, whereas informal channels enable breadth and speed.

This dual diffusion may have far-reaching implications for productivity, inequality, and the governance of innovation systems. Informal diffusion allows rapid experimentation and wide reach but often lacks the complementary investments in skills, workflow integration, and supervision necessary to generate strong productivity gains. Formal diffusion, in turn, embeds AI more deeply and productively but remains limited to firms that can bear the fixed costs of training and coordination. This asymmetry helps explain why widespread AI use among workers coexists with modest aggregate productivity improvements – a continuation of the modern productivity paradox (Brynjolfsson et al., 2017; Schweikl and Obermaier, 2020). Moreover, because informal diffusion depends on self-selection, it risks amplifying inequalities by education, gender, and age (Humlum and Vestergaard, 2025, 2024; Otis et al., 2024). Formal adoption, by contrast, can foster inclusion when paired with structured training and skill development, but its selective rollout confines benefits to specific segments of the workforce. Consequently, innovation systems built for firm-led technological change may overlook large portions of bottom-up experimentation, calling for new governance approaches that integrate both channels (Colovic et al., 2025; Fagerberg et al., 2005).

Synthesising these insights, we derive three guiding research questions. First, how widespread is AI use among workers, and how does its intensity vary by the formality of adoption? Second, what individual, establishment, and regional factors shape AI usage, and do these differ between formal and informal channels? Third, how does the formality of adoption relate to intensity of usage, training, AI-based supervision, and perceived productivity? These questions bridge diffusion theory with empirical evidence on the determinants and implications of AI use.

3 Data & Descriptive Statistics

3.1 The DiWaBe 2.0 Survey

The *Digital Transformation and the Changing World of Work (DiWaBe 2.0) survey* is an employee survey conducted between July and December 2024. The survey comprises around 9,800 respondents and is designed to be representative of the German workforce. It provides detailed information on the employees as well as their workplace, most notably on the use of artificial intelligence (AI) and digital technologies, tasks and skill requirements, organisational structures, training participation, and self-reported health.

Survey objective and structure. The goal of the DiWaBe 2.0 survey is to provide a representative picture of technological change at the workplace level in Germany, with a particular focus on the diffusion and use of artificial intelligence (AI). The DiWaBe 2.0 is part of a linked

employer-employee dataset that builds on a previous firm survey (BIZA II) in 2021.¹ Conducted in 2024, this survey constitutes the second wave of the DiWaBe study, following the initial wave in 2019. Together, both waves form a repeated cross-section with a smaller panel component that enables longitudinal analyses of technological change and its consequences for workers.

The target population includes all employees in Germany subject to social security contributions² who were employed at one of the BIZA II establishments as of June 30, 2021. The sample covers individuals aged 16 to 65 years and was drawn using a stratified random design based on qualification level, firm size, and age group. With the application of survey weights (see Section 3.1 for details), the survey provides a representative snapshot of all socially insured employees in Germany as of mid-2021.

The survey followed a mixed-mode design, combining Computer-Assisted Web Interviews (CAWI) and Computer-Assisted Telephone Interviews (CATI). Contact information was obtained from existing IAB records and earlier survey waves: 62.9 percent of sampled individuals ($N = 61,677$) included a phone number, while 37.1 percent ($N = 36,372$) did not. The survey was fielded by the SOKO Institute, which also supplemented missing contact information using public directories. After an extensive pretest, the questionnaire was streamlined to reduce length and respondent burden. The main field phase ran from July 18 to December 1, 2024.

In total, 9,835 interviews were completed, of which 9,410 were fully usable. Most responses were collected online via CAWI (76.2 percent), while 23.8 percent were obtained via CATI. The average interview lasted 34 minutes, with web interviews averaging 32 minutes and telephone interviews 41 minutes. Each interview was conducted independently of the employer, and all respondents were informed about the voluntary nature of participation and strict data protection standards.

Content. The DiWaBe 2.0 questionnaire covers a broad range of topics related to technological change and work in Germany. Its primary focus is on the diffusion and use of artificial intelligence (AI) at the workplace, but it also contains detailed information on work tasks, training activities, working conditions, and self-assessed health and well-being.

AI use is defined in the survey as "enabling computer programs and machines to independently perform tasks that would normally require human intelligence."³ The survey captures the diffusion of AI along both the extensive and intensive margins of adoption. At the extensive margin, respondents indicate whether they use any AI-based software or tool in their job (*AI users*) or not (*non-AI users*). At the intensive margin, they report the frequency of AI use on a five-point Likert scale ranging from "never" to "always."

Respondents were asked how often they use AI supported tools across a set of specific application types, including tools that (1) process text, (2) spoken language, (3) images or videos, (4) generate diagnostic or analytical outputs, (5) interact physically with humans (for example collaborative robots), (6) or fall into other categories. The reported frequencies at the application level are then aggregated to construct an overall measure of AI use, capturing both the range of applications used and the intensity of use.

In addition, respondents reported the single most important AI application they use at work in an open-text question. Based on this information, we classify AI users into those

¹The BIZA II establishment survey collected information on technology use at the firm level, focusing primarily on investment responses during the COVID-19 pandemic rather than on AI adoption. Accordingly, the present analysis relies on the employee survey component.

²Thus, apprentices, marginally employed workers, and student employees are excluded.

³Prior to this definition, respondents were asked whether they felt able to explain what artificial intelligence means. 79 percent indicated to be at least to some extend able to do so, while only around 7 percent indicated that they are not at all familiar with the concept of AI. Nevertheless, after assessing the familiarity, the term was explicitly defined to ensure a consistent understanding throughout the survey and across respondents.

whose main application involves generative AI and those relying primarily on non-generative AI technologies. Respondents were further asked whether this main AI application was formally introduced by their employer, allowing us to distinguish between *formal* and *non-formal* AI user. In our terminology, formality thus refers to whether the AI tool that respondents consider most important was officially provided by their employer. Thus non-formal users employ tools that have not been formally provided but are used on their own initiative. Importantly, non formal users may still be exposed to other AI applications that were formally introduced in their workplace.

Furthermore, respondents were asked to compare their AI use in 2024 with their use in 2019, enabling us to distinguish between workers with stable use and those who have increased their use over the past five years. This information provides a dynamic perspective on AI diffusion over time.

Beyond technological adoption, the survey collects detailed information on job tasks, skill requirements, and training. Respondents report how frequently they perform various types of tasks, including writing, math, programming, and manual or interpersonal work, as well as the degree of autonomy and time pressure they experience. It further covers workers' training uptake: whether respondents participated in any training during the past twelve months, the number of training days, whether training occurred during working hours, and whether it focused on job-specific, interdisciplinary, or digital skills. Additional questions capture how often employees use digital tools as part of their learning activities, providing insights into technology-supported training.

Finally, the survey allows to analyse perceived effects of technological change on work outcomes and well-being. Respondents assess, how their self-perceived productivity changes in response to AI use (regarding quality and quantity improvements as well as time saving), how satisfied they are with their work situation as a whole, as well as health-related indicators such as stress, exhaustion, and psychosomatic complaints.

Taken together, this broad range of questions makes the DiWaBe 2.0 survey a rich source for studying the interplay between AI use, task composition, skill development, and worker well-being.

Sampling, weighting, and representativeness. The DiWaBe 2.0 survey was drawn as a stratified random sample of socially insured employees who were employed at one of the BIZA II establishments as of June 30, 2021. Stratification was based on three key dimensions: skill level, firm size, and age group. Sampling probabilities therefore varied across strata to ensure adequate coverage of smaller subgroups of the workforce. The field sample comprised 98,049 individuals, of whom 9,835 ultimately participated in the survey. Within the realized sample, 9,410 interviews were fully usable for analysis.

To correct for unequal selection probabilities and selective non-response, a two-stage weighting procedure was applied. In the first stage, individual base weights were computed by combining the firm-level design weights from BIZA II with the employee-level sampling probabilities within each stratum. In the second stage, post-stratification weights were derived so that the weighted sample matched the joint distribution of skill, age, and firm-size categories in the underlying population. Weights were trimmed at the 99th percentile to prevent single observations biasing the results. A detailed documentation of the sampling and weighting procedures is provided in (Arntz et al., 2025).

After weighting, the DiWaBe 2.0 data reflects the structure of the socially insured workforce. As shown in Table A.1, the weighted DiWaBe 2.0 data mirrors the official employment statistics almost exactly across the three stratification dimensions. For example, medium-skilled employees represent 49.4 percent of the target population and 52.8 percent of the weighted sample,

while younger, middle-aged, and older workers are represented in virtually identical proportions to the population. Similarly, the distribution across firm-size classes aligns closely with administrative records.

Summary statistics. Table 1 provides an overview of baseline individual-, establishment-, and regional-level characteristics across AI user groups and highlights systematic compositional differences. Compared to non-users, both informal and formal AI users are slightly younger and less likely to be female, and they are substantially more educated. The share without any qualification is about 14 percent among non-users but only 6 percent among both user groups, while the share holding a college or university degree is roughly twice as high among users (33 and 36 percent) as among non-users (19 percent). These educational differences are reflected in job characteristics. AI users are much more likely to work in complex jobs (57 and 60 percent versus 41 percent) and less likely to be in skilled jobs (31 and 34 percent versus 45 percent). They are also more concentrated in business and administrative as well as IT and science-related occupations, whereas non-users are overrepresented in personal services. Differences at the establishment level are smaller but consistent with this pattern. Users tend to work in establishments with a more educated and more complex workforce, and formal AI users are more frequently employed in knowledge-intensive tertiary and ICT establishments and less often in public service. Regional characteristics vary little across groups, with only modestly higher population density and academic shares in regions where AI users reside. While these patterns are descriptive, they provide useful context for the regression analysis that follows, which evaluates which baseline (2019) characteristics are statistically and economically associated with subsequent informal and formal AI use.

Table 1: Summary Statistics by AI User Group


	Individual-level variables					
	Non-Users		AI Users (Non-Formal)		AI Users (Formal)	
	Mean	SD	Mean	SD	Mean	SD
Age (years)	48.15	11.05	45.79	11.39	45.82	11.51
Female	0.56	0.50	0.51	0.50	0.48	0.50
No educational qualification	0.14	0.35	0.06	0.24	0.06	0.24
Completed vocational training	0.61	0.49	0.46	0.50	0.47	0.50
Master craftsman/technician	0.06	0.23	0.12	0.33	0.13	0.34
College/University Degree	0.19	0.39	0.36	0.48	0.33	0.47
Skilled Job	0.45	0.50	0.31	0.46	0.34	0.47
Complex Job	0.41	0.49	0.60	0.49	0.57	0.50
Occupational Labor Market Tightness	1.27	1.41	1.25	1.39	1.19	1.37
Production Occupations	0.23	0.42	0.19	0.39	0.18	0.39
Personal Service Occupations	0.36	0.48	0.30	0.46	0.20	0.40
Business and Administrative Occupations	0.28	0.45	0.37	0.48	0.44	0.50
IT and Science-related Occupations	0.04	0.19	0.10	0.31	0.12	0.33
Other Business Services Occupations	0.08	0.27	0.04	0.19	0.06	0.23
Establishment-level variables						
	Non-Users		AI Users (Non-Formal)		AI Users (Formal)	
	Mean	SD	Mean	SD	Mean	SD
	43.61	4.71	43.22	5.29	42.17	5.20
Mean workforce age (years)	0.54	0.27	0.50	0.25	0.47	0.25
Workforce share female	0.69	0.21	0.60	0.24	0.59	0.25
Workforce share vocational educated	0.21	0.20	0.30	0.23	0.30	0.25
Workforce share college educated	0.57	0.23	0.52	0.26	0.49	0.26
Workforce share skilled	0.27	0.24	0.37	0.27	0.39	0.29
Workforce share complex	87.33	196.40	93.61	197.84	122.03	239.90
Establishment Labor Market Tightness	0.35	0.48	0.32	0.47	0.21	0.41
Public Service	0.07	0.25	0.10	0.30	0.07	0.25
Est. Size: Below 10 Employees	0.29	0.46	0.30	0.46	0.27	0.45
Est. Size: 10-49 Employees	0.33	0.47	0.28	0.45	0.30	0.46
Est. Size: 50-249 Employees	0.31	0.46	0.32	0.47	0.36	0.48
BIZA: Secondary, Knowledge-Intensive	0.08	0.26	0.08	0.27	0.09	0.28
BIZA: Secondary, Non-Knowledge-Intensive	0.19	0.39	0.14	0.34	0.14	0.34
BIZA: Tertiary, Knowledge-Intensive	0.23	0.42	0.24	0.43	0.31	0.46
BIZA: Tertiary, Non-Knowledge-Intensive	0.47	0.50	0.45	0.50	0.35	0.48
BIZA: ICT Sector	0.04	0.18	0.10	0.29	0.12	0.32
Regional variables						
	Non-Users		AI Users (Non-Formal)		AI Users (Formal)	
	Mean	SD	Mean	SD	Mean	SD
	44.25	2.11	43.75	2.07	43.75	2.10
Mean age	50.60	0.65	50.59	0.71	50.60	0.67
Share female (%)	15.11	6.75	17.06	7.77	17.00	7.86
Share academic (%)	0.45	0.22	0.45	0.23	0.42	0.21
Regional Labor Market Tightness	829.65	1020.55	1032.06	1146.61	1016.41	1218.74
Population density						

Notes: This table reports group-specific means and standard deviations for baseline individual characteristics, linked establishment characteristics, and regional characteristics, separately for non-users, informal AI users, and formal AI users. Individuals are classified as AI users if they report using AI-based tools at least occasionally in the 2024 DiWaBe 2.0 survey; among users, *formal* (employer-led) adoption indicates that the respondent reports their *main* AI tool was introduced by the employer, whereas *non-formal* (employee-led) adoption indicates that the main tool was not employer-introduced. All characteristics refer to pre-AI baseline information from 2019 (or the closest year to 2019 as possible, for the small minority of cases with missing values in 2019). Indicator variables are coded as 0/1 and reported as shares. Occupational, regional, and establishment labor market tightness is obtained from Bossler and Popp (2024). "Skilled job" and "complex job" is defined by the 5th digit of the KldB-2010 classification. Observations are weighted as described in Section 3.1.

3.2 The Distribution of AI at the Workplace

Before turning to the multivariate analyses, we first describe how frequently AI is used in the workplace and how usage differs by application and mode of introduction. We distinguish between the usage at the extensive (if AI is used) and intensive margins (how often AI is used), variation across application types, and differences depending on whether AI use is formally introduced by the employer or informally initiated by the employee. These patterns provide an initial empirical foundation for the analysis of AI diffusion mechanisms and their implications later on.

Figure 1: Frequency of AI Usage (Overall)



Note: This figure shows the weighted percentage distribution of respondents by how frequently they use AI-supported computer programs at the workplace. Response categories are aggregated to *Never*, *Rarely/Sometimes*, and *Often/Always*. Observations are weighted using trimmed post-stratification weights described in Section 3.1.


Source: DiWaBe 2.0 Survey

Frequency and formality of AI use. Figure 1 shows that 64 percent of employees report using an AI-based application at least occasionally in 2024 and are therefore classified as *AI users*. However, only about 20 percent report regular use (that is, “often” or “always”), indicating that while AI tools have reached a majority of workplaces, they are not yet deeply embedded in most workers’ day-to-day routines. This high prevalence at the worker level contrasts sharply with adoption rates reported in firm surveys, which typically range from about 10 to 27 percent (e.g., Gerhards and Baum, 2024; Licht and Wohlrabe, 2024; Falck et al., 2024).

Indeed, we observe similarly low rates of firm-led AI adoption in our data. Only about one third of AI users report that their main AI application was formally introduced by their employer. By contrast, the majority rely on tools adopted informally, potentially without explicit organizational support, integration, or even the employer’s knowledge. One likely explanation is the broad accessibility of many generative AI tools, which are often freely available, require little or no organizational integration, and can be used across both professional and personal contexts.

The frequency of AI use differs sharply between formal and informal users. As shown in Figure 2 (top panel), nearly half of formal users report regular use (“often” or “always”), compared with only about one quarter of non-formal users. Among informal users, occasional or rare use clearly dominates. To assess whether formal introduction is associated with deeper integration of AI beyond usage itself, we later examine differences in workplace routines, training, supervision, and productivity.

Figure 2: Frequency of AI Usage by Diffusion Mode



Note: This figure shows the weighted percentage distribution of respondents by frequency of AI use, separately for AI users whose main AI tool was not formally introduced by the employer (informal) and those whose main AI tool was introduced by the employer (formal). The upper panel reports overall usage frequency by formality. The lower panel reports frequency by AI application category (text, spoken language, images/videos, diagnostic/analytical tools, collaborative systems, and other). Frequency responses are aggregated into *Never*, *Rarely/Sometimes*, and *Often/Always*. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey

Application types and entry barriers. The type of AI tools used also varies by formality. Figure 2 (bottom panel) shows usage across six AI application categories. Text-based generative AI (e.g., ChatGPT) is by far the most common, used by over 80 percent of all AI users, regardless of formality – reflecting low entry barriers and high accessibility. However, formal users report more regular use even for these tools (27 percent vs. 18 percent).

Differences by formality become more pronounced for applications with higher adoption costs or greater integration requirements (e.g., legal concerns). The most pronounced gap is with diagnostic tools, which are used by around 50 percent of formal users but only 30 percent of informal users. More costly AI software that processes spoken language or images and videos also shows a gap in adoption of 7 percent and 9 percent respectively. The same holds true for collaborative AI systems such as robots with a 10 percent difference in usage probability. These patterns suggest that informal adoption is more concentrated in accessible, low-cost tools, whereas more complex and specialized technologies tend to diffuse through employer-led channels.

This observation highlights an important distinction between AI and previous technological

advances. Whereas earlier waves of technological change typically required substantial firm-level investments and were "top down," many AI tools enable individual experimentation at minimal cost, leading to "bottom up" technological change. On the one hand, this easy access to AI may drive its rapid diffusion. At the same time, informal adoption depends on self-selection which is driven by factors such as motivation, technological openness, or peer effects – characteristics that tend to vary systematically by gender, age, or education; which may reinforce existing inequalities among these groups. Formal implementation, by contrast, can create broad opportunities for structured training and more equitable access. On the other hand it means standardization and less flexibility. Thus, understanding who adopts AI – and through which pathway – is therefore critical for assessing how modern AI technologies reshape workplaces and create or reduce economic inequality.

4 Empirical Design

Our primary goal is to describe how AI users differ from non-users and how these relationships vary depending on whether AI was formally introduced by the employer or informally initiated by the employee. First, to examine how pre-existing individual, establishment, and regional characteristics are associated with the (increasing) use of artificial intelligence (AI) at the workplace, we estimate a series of linear probability models (LPMs). Then we formally test if these patterns differ by formality, and if formality is correlated with a different intensity of usage, continuing training, supervision, and perceived productivity.

4.1 The Determinants of AI Usage

To examine which pre-AI characteristics are associated with the diffusion of AI and with usage at the extensive and intensive margins, we estimate the following regression model:

$$y_i = \alpha + \mathbf{X}_i' \boldsymbol{\beta} + \mathbf{F}_{j(i)'} \boldsymbol{\gamma} + \mathbf{R}_{r(i)'} \boldsymbol{\delta} + \varepsilon_i. \quad (1)$$

The outcome variable y_i captures three AI-related measures: (i) whether the employee uses AI at work in 2024, (ii) whether AI use at work increased between 2019 and 2024, and (iii) the intensity of AI use in 2024, measured on a five-point Likert scale. A detailed description of these outcomes is provided in Panel A of Table 2. For robustness, we re-estimate the intensity models using ordered logit specifications and find qualitatively similar results.

Explanatory variables. \mathbf{X}_i is a vector of individual-level characteristics, $\mathbf{F}_{j(i)}$ reflects establishment-level characteristics for establishment j in which individual i is employed, and $\mathbf{R}_{r(i)}$ captures regional controls at the county level (more details below). At the individual level, prior research identifies age, gender, education, and personality traits as key predictors of AI adoption (Bick et al., 2024; Humlum and Vestergaard, 2024; Otis et al., 2024; Kaya et al., 2024). We therefore include controls for gender, age, and educational attainment, distinguishing between no vocational degree, vocational degree, master craftsman or technician, and college degree, as well as indicators for job-skill complexity (low-skilled, skilled, or complex). To capture attitudinal and cognitive dimensions related to technology adoption, we control for personality traits using the Big Five framework, where openness is closely linked to experimentation with new technologies (Kaya et al., 2024). Finally, we include occupational labor market tightness, measured as the occupation-specific vacancy-to-unemployment ratio in the local labor market (Bossler and Popp, 2024).

At the establishment level, prior studies highlight organizational resources, learning environments, and sectoral characteristics as important drivers of formal AI diffusion (Falck et al.,

2024; Na et al., 2023; Gualandri and Kuzior, 2024; Gerhards and Baum, 2024). Accordingly, we control for public-sector status, average workforce age, the share of female employees, and the distribution of employees across education and skill levels in the establishment, mirroring the individual-level categories. These variables proxy the establishment's absorptive capacity and its ability to invest in complementary inputs such as training or task redesign. We additionally include establishment-level labor market tightness, defined as the weighted average of occupational tightness across the establishment's employment structure.

Finally, regional diffusion processes depend on local capabilities and network spillovers that shape both employer- and employee-driven adoption (Dahlke et al., 2024; Na et al., 2023; Nicoletti et al., 2020). To account for these factors, we include regional controls for demographic structure, educational attainment, population density, and labor market tightness (Bossler and Popp, 2024). This multilevel framework allows us to assess how individual, organizational, and regional factors jointly influence AI diffusion and its mode of adoption.

Robustness and interpretation Importantly, all covariates are measured in 2019 – prior to the widespread diffusion of generative AI. This ensures that all explanatory variables temporally precede the outcome, mitigating concerns regarding simultaneity. To account for potential unobserved factors (captured in the error term ε_i) among individuals working at the same establishment – such as work culture, shared exposure to organizational change, or AI-related training policies – we cluster standard errors at the establishment level. To further absorb differences in AI adoption that reflect sector- or occupation-specific trends – such as structural change, regulatory environments, or task structures that influence the adoption rate of AI – we augment our baseline specification by including fixed effects for occupation segment and industry:

$$y_i = \alpha + \mathbf{X}'_i \boldsymbol{\beta} + \mathbf{F}'_{j(i)} \boldsymbol{\gamma} + \mathbf{R}'_{r(i)} \boldsymbol{\delta} + \rho_{o(i)} + \zeta_{s(i)} + \varepsilon_i \quad (2)$$

where $\rho_{o(i)}$ denotes fixed effects for the respondent's 2-digit occupation and $\zeta_{s(i)}$ denotes fixed effects for the 2-digit industry classification of their establishment. The inclusion of these fixed effects ensures that comparisons are drawn within rather than across occupational and sectoral contexts, reducing potential bias from omitted group-level characteristics that may confound the interpretation of observed associations.

4.2 Differences by Formality Among AI Users

Next, we examine differences between formal and informal AI users in terms of usage intensity, application-specific usage patterns, training participation, AI-based supervision, and perceived productivity gains (Figure 5). We restrict the sample to individuals who report using AI and estimate the following specification:

$$y_i = \alpha + \lambda, \text{Formal}_i + \mathbf{X}'_i \boldsymbol{\beta} + \mathbf{F}'_{j(i)} \boldsymbol{\gamma} + \mathbf{R}'_{r(i)} \boldsymbol{\delta} + \rho_{o(i)} + \zeta_{s(i)} + \varepsilon_i, \quad (3)$$

where y_i captures five outcome dimensions among AI users (definitions, means, and standard deviations are reported in Panels B to D of Table 2):

- **Overall intensity of AI usage:** a 4-point Likert scale ranging from rarely to always.
- **Application-specific usage intensity:** indicators for intensive AI use in text processing, voice processing, image processing, diagnosis creation, coding, and other applications.
- **Perceived productivity gains:** 5-point Likert scales measuring self-assessed improvements in time or speed, quantity, and quality of work due to AI use.

- **AI-based supervision (algorithmic management):** binary indicators for whether task allocation, time management, or performance management is supervised by AI.
- **Training participation:** a binary indicator for participation in any job-related continuing training in the past twelve months, and the number of distinct training courses attended.

The variable Formal_i equals one if the respondent reports that their main AI application was formally introduced by the employer, and zero otherwise. All control variables and fixed effects are identical to those used in Equation (2). The coefficient λ therefore captures differences between formal and informal AI users within the same occupation and industry, conditional on individual, establishment, and regional characteristics.

Table 2: Definition and Summary Statistics of Main Outcome Variables

Variable	Description (English Translation)	Mean	SD
A. Usage			
Any AI use	Uses at least one AI tool at work (0/1)	0.644	0.479
Intensity of AI use	Intensity of AI use (0 = never, 4 = always)	1.288	1.199
Increase in AI use since 2019	Increase in AI use since 2019 (0/1)	0.267	0.442
B. Perceived Productivity			
Quantity improvement	“With the help of AI, tasks are taken off my hands.” (0–4)	1.948	1.270
Quality improvement	“With the help of AI, I achieve higher-quality work results.” (0–4)	2.053	1.205
Time improvement	“With the help of AI, I work faster or produce more.” (0–4)	2.274	1.311
C. Supervision / Algorithmic Management			
Time management	Working hours regularly determined by AI (0/1)	0.025	0.157
Task allocation	Tasks regularly assigned by AI (0/1)	0.033	0.180
Performance management	Work performance evaluated by AI (0/1)	0.009	0.097
D. Training			
Training participation	Participated in continuing training (0/1)	0.444	0.497
Number of courses	Number of continuing training courses attended	1.651	3.580

Notes: This table lists main outcome variables, a short description, and weighted means and standard deviations. Indicator variables are coded as 0/1 and reported as shares. Questions regarding perceived productivity are only asked to AI users. Observations are weighted as described in Section 3.1.

Interpretation. We refrain from interpreting estimated coefficients as causal effects, but rather see them as meaningful associations between (i) pre-existing individual characteristics, workplace structures, and local labor market conditions and the diffusion of AI, and (ii) the formality of AI use – “bottom-up” by the worker (informal usage) versus “top-down” by the employer (formal usage) – and subsequent workplace practices. Together, these models provide

a comprehensive view of how AI diffuses across different levels of the labor market and how the mode of diffusion conditions its intensity, training incidence, supervision patterns, and perceived productivity gains. The next section presents the empirical results and discusses how these associations inform our understanding of dual diffusion in the age of AI.

5 Results


5.1 Patterns in the Usage and Diffusion of AI

We begin by examining which pre-existing individual, establishment, and regional characteristics are associated with the adoption of AI. Figure 3 presents the estimated coefficients from a joint weighted linear probability model (LPM) predicting whether employees increasingly use AI at work in 2024 compared to 2019 – either without additional fixed effects, with 2-digit industry fixed effects, and with 2-digit industry and 2-digit occupation fixed effects (for details, see Section 4). The results align closely with prior evidence that education is a key predictor of digitalization and AI adoption (Bick et al., 2024; Humlum and Vestergaard, 2024) as employees with a college degree or a master craftsperson/technician qualification are 10 to 15 p.p. more likely to have increased their use of AI. Given that roughly 65 percent of employees report using AI at least occasionally, these differences are economically substantial and reinforce the skill-biased nature of technological diffusion observed in earlier studies. Women are around 6 p.p. less likely to report an increase in AI use – a pattern also observed in early studies on generative AI adoption (Otis et al., 2024; Chugunova et al., 2026). Among the Big Five traits, openness and extraversion show a small (2.5 p.p.) but consistent positive association with increasing AI use, supporting prior evidence that individuals open to new experiences are more willing to experiment with new technologies (Kaya et al., 2024). Other traits and occupational labor market tightness show no systematic effects.

In contrast to individual characteristics, establishment and regional factors play only a minor role. The only exception is the public sector, which is significantly negatively associated with a reduced adoption of AI by around 7 p.p. Establishments in the public sector and with a high share of vocationally trained workers are both associated with a lower likelihood of AI use, though these effects lose significance once fixed effects are introduced. Mean workforce age (weakly positive) and the share of female employees (weakly negative) show only a weak correlation, while regional characteristics do not have a significant relationship throughout.

Figure A.2 examines the characteristics associated with any use of AI at work in 2024. Reassuringly, the same pattern emerges: Master craftspersons or technicians have the highest probability of using AI – about 35 p.p. higher than employees without a qualification – followed by those with a college or university degree (+25 p.p.). While the relationship with gender is less pronounced, age shows a negative relationship (-4 p.p. per standard-deviation increase), consistent with findings that younger workers adopt new technologies more readily (Humlum and Vestergaard, 2024).

Figure 3: Determinants of Increasing AI Usage



Note: This figure shows the results from a weighted linear probability model regressing whether respondents report an increase in AI use between 2019 and 2024 on a set of individual, establishment, and regional variables. Covariates are measured in 2019 (pre-generative-AI diffusion). The figure compares specifications with (i) no fixed effects, (ii) 2-digit industry fixed effects, and (iii) 2-digit industry and 2-digit occupation fixed effects (see Equation 2 in Section 4). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.


Source: DiWaBe 2.0 Survey

The results for the intensive margin in Appendix Figure A.1 mirror those for the extensive margin. Master craftspersons and technicians report the highest intensity of AI use, with levels about 0.7 Likert points higher than employees without any qualification, followed by college graduates (+0.45 points). Given that the average intensity of AI use is 1.3 points (see Table 2), these differences are substantial and emphasize that education is the central determinant of AI diffusion. Again, gender does not show a significant negative relationship with usage inten-

sity once education and occupation are accounted for, while age does. Similarly, public-sector establishments and those with a high share of vocationally trained workers use AI somewhat less intensively, although this difference becomes statistically insignificant once we control for occupation and industry differences.

5.2 The Role of Formality

Figure 4: Determinants by Formality



Note: This figure shows the results from a weighted linear probability model regressing whether respondents report an (i) formal or (ii) informal increase in AI use between 2019 and 2024 on a set of individual, establishment, and regional variables, as well as 2-digit industry and 2-digit occupation fixed effects (see Equation 2 in Section 4). Covariates are measured in 2019 (pre-generative-AI diffusion). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey

These results indicate that AI diffusion at the workplace is primarily driven by bottom-up adoption, shaped by individual skills, job complexity, and willingness to experiment rather than by coordinated organizational strategies. This pattern implies that early AI use is positively selected, raising concerns that worker-initiated diffusion may reinforce existing inequalities in access, learning opportunities, and returns to technology. A key question, therefore, is whether employer-led AI implementation alters this dynamic. In the following section, we examine whether formal adoption broadens access to AI use or instead primarily affects how AI is embedded in organizations. Specifically, we assess whether formality is associated with deeper integration through higher usage intensity, structured training, AI-based supervision, and perceived productivity gains.


The characteristics associated with formal and informal increases in AI use. Figure 4 compares the determinants of an increase in AI use between 2019 and 2024, distinguishing between increases driven by formally introduced tools and those arising from informal, self-initiated use. The patterns are remarkably similar across both diffusion modes. Increases in AI use, whether formal or informal, are strongly associated with higher educational attainment and more complex jobs, indicating that intensification remains skill-biased regardless of how AI enters the workplace. This suggests that formal introduction does not primarily redirect AI uptake toward different worker groups, but rather operates on a similar pool of positively selected employees.

Differences across establishment and regional characteristics are limited. The main exception is the public sector, where employees are significantly less likely to experience a formal increase in AI use. This pattern is consistent with prior evidence showing that institutional constraints, regulatory requirements, and limited organizational readiness (and knowledge) slow down employer-led AI implementation (Falck et al., 2024). Taken together, these results indicate that employer-led introduction has not fundamentally altered who increases their use of AI. Instead, formality appears to shape how AI use evolves conditional on adoption, rather than broadening the set of workers who intensify their use.

Formality and the implementation of AI. We now turn to how the mode of introduction relates to the embedding of AI at the workplace. This analysis addresses two related questions. First, whether formal introduction is associated with different implementation practices, such as more intensive use, structured training, or algorithmic coordination. Second, whether formality is linked to higher productivity gains from AI use, potentially helping to explain why recent studies find limited aggregate productivity effects despite widespread adoption (Acemoglu et al., 2022; Humlum and Vestergaard, 2025).

As Figure 1 in Section 3.2 already indicated, formal and informal users differ in how they use AI: formal users apply AI more intensively and for more specific, task-related purposes, while informal users tend to rely more on general, low-cost text applications. Figure 5 shows that these differences persist in multivariate regressions that condition on individual, establishment, and regional characteristics as well as industry and occupation fixed effects. Formal introduction is associated with significantly higher overall usage intensity, consistent with evidence that organizational integration fosters more intensive use of new technologies (Na et al., 2023; Gualandri and Kuzior, 2024).

Figure 5: Formality and Associated Factors



Note: This figure shows associations between formal AI introduction and AI usage intensity, AI application type, perceived productivity, AI-based supervision, and training participation, conditional on being an AI user. Estimates are based on weighted linear probability models and control for individual, establishment, and regional characteristics, as well as 2-digit industry and 2-digit occupation fixed effects (see Equation 3 in Section 4). Covariates are measured in 2019 (pre-generative-AI diffusion). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey

In line with this interpretation, formality is also associated with a broader and more advanced application profile. Differences are small for low-cost, easily accessible tools such as text and voice processing, but substantially larger for applications that require organizational coordination or investment, including image processing, diagnostic or analytical systems, and collaborative technologies such as cobots. This pattern suggests that employer-led adoption primarily facilitates the use of AI where complementary infrastructure, legal clearance, or workflow integration are required.

The introduction of new technologies by employers is typically accompanied by new skill requirements and task adjustments (Arntz et al., 2016; Acemoglu and Autor, 2011). Continuing training is therefore a central mechanism through which firms adapt to technological change

(Muehlemann, 2025; Heß et al., 2023). Figure 5 shows that, conditional on the full set of controls, workers whose employer formally introduced their main AI tool are around 8 p.p. more likely to participate in training and attend, on average, about 0.4 additional courses compared to informal users. This association indicates that formal AI adoption is systematically coupled with complementary investments in human capital, consistent with deeper organizational learning rather than a mere expansion of access.

A further dimension of AI implementation concerns monitoring and coordination. Formal AI adoption is positively associated with the use of AI systems for task allocation, time management, and performance evaluation (Figure 5). Although the overall prevalence of AI-based supervision remains low, the relative differences are economically meaningful: formal users are 3 to 7 p.p. more likely to report algorithmic management, compared to baseline rates of 2.5, 3.3, and 0.9 percent, respectively (Table 2).

To assess whether these patterns reflect organizational implementation rather than individual behavior, we re-estimate the models in the full sample and distinguish between formal and informal AI use. If AI-based supervision primarily reflects employer-driven adoption, it should be concentrated among formal users, whereas informal users should resemble non-users. Table A.2 confirms this prediction: the association with AI-based supervision is largely driven by formal AI use, while informal users differ only marginally from non-users. Expressed in standard deviation units, supervision rates are 0.33 to 0.50 SD higher among formal users, compared to 0.05 to 0.14 SD for informal users. While this pattern is consistent with greater technological sophistication and organizational capacity, it also highlights potential trade-offs between efficiency gains and increased monitoring intensity.

Finally, we examine whether formal AI adoption is associated with higher perceived productivity gains. Figure 5 shows that formal users report significantly stronger improvements in work quality, quantity, and time efficiency. Estimated differences range from 0.39 to 0.48 Likert points relative to baseline means of 1.95 to 2.27 (Table 2). These associations remain robust when additionally controlling for usage intensity, application types, training participation, and AI-based supervision (Appendix Table A.3). Overall, these patterns suggest that employer-led AI adoption is associated with higher perceived productivity improvements – even beyond associated factors such as task supervision, higher usage intensity, different application types, and the provision of training.

As an additional robustness check, we include establishment fixed effects to absorb time-invariant unobserved heterogeneity across workplaces that could jointly shape AI implementation and perceived productivity. This approach compares AI-using employees within the same establishment who differ in whether their main AI application was introduced by the employer or adopted on their own initiative. As shown in Figure A.3, differences in training participation, the number of training courses attended, and AI-based supervision are no longer statistically distinguishable from zero once establishment fixed effects are included. Similarly, differences in low-cost application types such as text and voice processing disappear. In contrast, differences in overall usage intensity and in more costly or infrastructure-intensive applications, including image processing, diagnostic tools, and cobots, persist within establishments. Importantly, perceived productivity gains also remain economically meaningful and statistically significant. Taken together, these results indicate that training and supervision differences primarily reflect between-establishment variation in implementation strategies and organization, whereas the association between formal introduction, intensive and advanced AI use, and higher perceived productivity operates to a large extend within establishments.

6 Conclusion

This paper provides representative worker-level evidence on the *dual diffusion* of artificial intelligence inside workplaces, distinguishing between formal, employer-led introduction and informal, employee-initiated use. Using linked employer-employee data for nearly 10,000 workers in Germany, we show that AI has diffused widely: almost two thirds of employees report using AI tools at least occasionally. At the same time, only around 20 percent of workers frequently use AI, indicating that broad diffusion has not (yet) translated into deep integration for most workers.

Across the extensive and intensive margins, diffusion is strongly skill-biased. Education and task complexity are strongly correlated with the adoption and usage of AI. Establishment and regional characteristics play a comparatively smaller role. This pattern is consistent with a technology diffusion process that is largely driven by individual initiative rather than by organizational strategy.

The distinction between formal (top-down) and informal (bottom-up) diffusion is crucial along two dimensions: diffusion and embedding. First, formal adoption of AI is less common as only one third of AI users report that their main AI-application has been introduced by their employer. At the same time, formality is associated with deeper workplace integration: formal users report more intensive use, more structured training, and a higher prevalence of AI-based coordination and supervision tools, alongside higher perceived productivity gains. However, formal introduction does not attenuate the skill bias of AI adoption. Formal and informal users are similarly positively selected by education and task complexity, implying that employer-led implementation has not (yet) broadened access to AI.

These findings help to interpret why widespread AI use can coexist with modest aggregate productivity gains. Informal diffusion enables rapid experimentation with low-entry tools, but it often occurs without complementary investments in training and workflow adaptation that are likely required for sustained productivity improvements. Formal diffusion is associated with precisely these complements, yet it remains less common, implying that deeper embedding – and the associated gains – may remain concentrated rather than broadly shared. Dual diffusion also raises distributional and governance challenges. Because AI currently spreads largely through self-initiated use, the same skill gradients that shape adoption can translate into divergence in learning and intensification over time. Consistent with this, women are less likely to have increased their AI use since 2019, and adoption and intensification are systematically lower in the public sector, pointing to institutional frictions. Moreover, bottom-up adoption may complicate coordination and compliance when external tools are used outside standardized processes, while formal embedding is linked to increased AI-based monitoring and supervision.

Looking ahead, a key challenge for research and policy is to understand how informal experimentation translates into formal organizational adoption over time and which complementary investments can turn broad exposure into sustained productivity gains. For organizational strategy and innovation governance, our results suggest a concrete implication: if top-down integration is to support inclusive diffusion, it needs to be designed not only to embed AI more deeply, but also to broaden access in order to not intensify inequalities.

References

Acemoglu, Daron and David Autor (2011) “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in Card, David and Orley Ashenfelter eds. *Handbook of Labor Economics*, 4, 1043–1171: Elsevier.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022) “Artificial intelligence and jobs: Evidence from online vacancies,” *Journal of Labor Economics*, 40 (S1), S293–S340.

Aldasoro, I., O. Armantier, S. Doerr, L. Gambacorta, and T. Oliviero (2024) “The gen AI gender gap,” *Economics Letters*, 241, 111814.

Arntz, Melanie, Myriam Baum, Eduard Brüll et al. (2025) “Digital Transformation and the Changing World of Work (DiWaBe 2.0): A Data Source for Research on Artificial Intelligence and Other Technologies in the Workplace,” Technical report, 10.21934/BAUA: BERICHT20250319.

Arntz, Melanie, Terry Gregory, Simon Jansen, and Ulrich Zierahn (2016) “Tätigkeitswandel und Weiterbildungsbedarf in der digitalen Transformation,” Research Report, ZEW-Gutachten und Forschungsberichte.

Battisti, Giuliana and Paul Stoneman (2003) “Inter- and intra-firm effects in the diffusion of new process technology,” *Research Policy*, 32 (9), 1641–1655.

——— (2010) “The Diffusion of New Technology,” in Hall, Bronwyn H. and Nathan Rosenberg eds. *Handbook of the Economics of Innovation*, 2, 733–760, Amsterdam: Elsevier.

Bick, Alexander, Adam Blandin, and David J. Deming (2024) “The Rapid Adoption of Generative AI,” September, 10.3386/w32966.

Bossler, Mario and Martin Popp (2024) “Labor Demand on a Tight Leash,” Technical report, IZA Discussion Papers.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025) “Generative AI at Work,” *The Quarterly Journal of Economics*, 140 (2), 889–942, 10.1093/qje/qjae044.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2017) “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics,” NBER Working Paper 24001, National Bureau of Economic Research.

Büchel, Jan and Roschan Monsef (2024) “Künstliche Intelligenz: Bessere Entlohnung durch Produktivitätsbooster?” *IW-Trends-Vierteljahresschrift zur empirischen Wirtschaftsforschung*, 51 (2), 45–63.

Chugunova, Marina, Dietmar Harhoff, Katharina Hözlle, Verena Kaschub, Sonal Malagimani, Ulrike Morgalla, and Robert Rose (2026) “Who uses AI in research, and for what? Large-scale survey evidence from Germany,” *Research Policy*, 55, 105381, 10.1016/j.respol.2025.105381, Short communication.

Cohen, Wesley M. and Daniel A. Levinthal (1990) “Absorptive Capacity: A New Perspective on Learning and Innovation,” *Administrative Science Quarterly*, 35 (1), 128–152.

Colovic, Ana, Annalisa Caloffi, Federica Rossi, and Margherita Russo (2025) “Institutionalising the digital transition: The role of digital innovation intermediaries,” *Research Policy*, 54 (1), 105146.

Dahlke, Johannes, Mathias Beck, Jan Kinne, David Lenz, Robert Dehghan, Martin Wörter, and Bernd Ebersberger (2024) “Epidemic effects in the diffusion of emerging digital technologies: evidence from artificial intelligence adoption,” *Research Policy*, 53 (2), 104917, <https://doi.org/10.1016/j.respol.2023.104917>.

Fagerberg, Jan, David C. Mowery, and Richard R. Nelson eds. (2005) *The Oxford Handbook of Innovation*, Oxford: Oxford University Press.

Falck, Oliver, Anna Kerkhof, and Anita Wölfel (2024) “Künstliche Intelligenz–wie Unternehmen sie nutzen und was sie noch daran hindert,” *ifo Schnelldienst*, 77 (09), 57–63.

Gerhards, Christian and Myriam Baum (2024) “AI in the Workplace: Who Is Using It and Why? A Look at the Driving Forces Behind Artificial Intelligence in German Companies,” *Annals of Computer Science and Information Systems*, 41, 45–52.

Giering, Oliver and Stefan Kirchner (2021) “Künstliche Intelligenz am Arbeitsplatz: Forschungsstand, Konzepte und empirische Zusammenhänge zu Autonomie,” *Soziale Welt*, 72 (4), 551–588, [10.5771/0038-6073-2021-4-551](https://doi.org/10.5771/0038-6073-2021-4-551).

Gualandri, Fabio and Aleksandra Kuzior (2024) “Drivers of AI adoption in enterprises: A European-wide analysis,” *Materials Research Proceedings*, 45.

Hall, Anja (2024) “IT skills: Requirements for employed persons,” *BWP* 53, 37–39.

Heß, Pascal, Simon Janssen, and Ute Leber (2023) “The Effect of Automation Technology on Workers’ Training Participation,” *Economics of Education Review*, 96, 102438, [10.1016/j.econedurev.2023.102438](https://doi.org/10.1016/j.econedurev.2023.102438).

von Hippel, Eric (2006) *Democratizing Innovation*, Cambridge, MA: MIT Press.

Humlum, Anders and Emilie Vestergaard (2024) “The Adoption of ChatGPT,” Working Paper 16992, IZA Discussion Papers.

——— (2025) “Large language models, small labor market effects,” Technical report, National Bureau of Economic Research.

Kaya, F., F. Aydin, A. Schepman, P. Rodway, O. Yetisensoy, and M. Demir Kaya (2024) “The roles of personality traits, AI anxiety, and demographic factors in attitudes toward artificial intelligence,” *International Journal of Human-Computer Interaction*, 40 (2), 497–514.

Kunze, Florian, Carolina Opitz, and Ann Sophie Lauterbach (2025) “Konstanzer KI-Studie 2025: Die Nutzung von Künstlicher Intelligenz in der Arbeitswelt steigt, Ungleichheiten in der Wahrnehmung bleiben weiterhin bestehen : Ergebnisbericht Juli 2025,” Technical report.

Licht, Thomas and Klaus Wohlrabe (2024) “Ai adoption among german firms,” *CESifo Working Paper*.

Lundvall, Bengt-Åke ed. (1992) *National Systems of Innovation: Towards a Theory of Innovation and Interactive Learning*, London: Pinter Publishers.

Muehlemann, Samuel (2025) “Artificial intelligence adoption and workplace training,” *Journal of Economic Behavior & Organization*, 238, 107206.

Na, Seunguk, Seokjae Heo, Wonjun Choi, Sehee Han, and Cheekyung Kim (2023) “Firm size and artificial intelligence (AI)-Based technology adoption: the role of corporate size in South Korean construction companies,” *Buildings*, 13 (4), 1066.

Nelson, Richard R. and Sidney G. Winter (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.

Nicoletti, Giuseppe, Christina von Rueden, and Dan Andrews (2020) “Digital technology diffusion: A matter of capabilities, incentives or both?” *European Economic Review*, 128, 103513, <https://doi.org/10.1016/j.eurocorev.2020.103513>.

Noy, Shakked and Whitney Zhang (2023) “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381 (6654), 187–192.

Otis, Nicholas, Yilin Chen, Michael Jones, and Shreya Raghunathan (2024) “Global Evidence on Gender Gaps and Generative AI,” HBS Working Paper 24-092, Harvard Business School, Harvard Business School Working Paper Series.

Polák, Petr (2017) “The Productivity Paradox: A Meta-Analysis,” *Information Economics and Policy*, 38, 38–54.

Rogers, Everett M. (2003) *Diffusion of Innovations*, New York: Free Press, 5th edition.

Schweikl, Stefan and Robert Obermaier (2020) “Lessons from Three Decades of IT Productivity Research: towards a better understanding of IT-induced productivity effects,” *Management Review Quarterly*, 70 (4), 461–507.

Van Noorden, Richard and Jeffrey M. Perkel (2023) “AI and science: what 1,600 researchers think,” *Nature*, 621, 672–675, 10.1038/d41586-023-02980-0, News Feature (published 27 Sep 2023; correction 10 Oct 2023).


A Appendix

Table A.1: Representativeness of the Survey Data Compared to the Target Population

	Target Population (1)	Survey Weighted (2)	Survey Unweighted (3)
<i>Stratification Variables (in Percent)</i>			
<i>Share of Individuals by Skill Level:</i>			
Lower-skilled	34.7	31.3	6.6
Medium-skilled	49.4	52.8	35.2
High-Skilled	15.9	15.8	58.2
<i>Share of Individuals by Age Group:</i>			
Younger than 35	28.0	28.0	30.7
35–49 years	35.2	35.2	36.8
50 or older	36.8	36.9	32.5
<i>Share of Individuals by Firm Size:</i>			
0–<50 employees	40.7	40.3	15.5
50–<200 employees	24.3	24.5	39.6
200 or more employees	35.0	35.2	45.0
Number of Individuals	29,672,942	9,835	9,835

Note: The table shows the distribution by selected stratification variables as of June 30, 2021, based on the IAB Employment History (BeH) V10.08.00-202112, Nuremberg 2023. The shares in column (2) are calculated using the post-stratification weight, while those in columns (1) and (3) are unweighted.


Figure A.1: Determinants of AI Usage Intensity



Note: This figure shows the results from a weighted linear probability model regressing whether respondents report the intensity of AI usage in 2024 (4-point Likert scale) on a set of individual, establishment, and regional variables. Covariates are measured in 2019 (pre-generative-AI diffusion). The figure compares specifications with (i) no fixed effects, (ii) 2-digit industry fixed effects, and (iii) 2-digit industry and 2-digit occupation fixed effects (see Equation 2 in Section 4). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey

Figure A.2: AI Usage Determinants



Note: This figure shows the results from a weighted linear probability model regressing whether respondents report the use of AI in 2024 on a set of individual, establishment, and regional variables. Covariates are measured in 2019 (pre-generative-AI diffusion). The figure compares specifications with (i) no fixed effects, (ii) 2-digit industry fixed effects, and (iii) 2-digit industry and 2-digit occupation fixed effects (see Equation 2 in Section 4). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey

Table A.2: User Type and Supervision

	AI time allocation		AI task allocation		AI performance measure	
	(1)	(2)	(3)	(4)	(5)	(6)
Informal user	0.0137 (0.00998)	0.0225** (0.0105)	0.0173** (0.00741)	0.0168** (0.00802)	0.00428 (0.00400)	0.00518 (0.00444)
Formal user	0.0601*** (0.0227)	0.0676*** (0.0225)	0.0850*** (0.0151)	0.0883*** (0.0139)	0.0308*** (0.00837)	0.0323*** (0.00843)
Mean of dep. var.	0.025	0.025	0.033	0.033	0.009	0.009
SD of dep. var.	0.157	0.157	0.180	0.180	0.097	0.097
Observations	7,102	7,100	6,819	6,816	7,019	7,016
R-squared	0.055	0.135	0.055	0.101	0.030	0.094
Additional controls	✓	✓	✓	✓	✓	✓
Fixed effects		✓		✓		✓


Notes: This figure shows the results from weighted linear probability models regressing whether respondents report the usage of AI-based management tools for (i) task allocation, (ii) time management, (iii) and performance evaluation on a categorical variable indicating the user type (non-user, informal AI user, formal AI user) while controlling for a set of individual, establishment, and regional variables. Columns (2), (4), and (6) additionally control for 2-digit industry and 2-digit occupation fixed effects (see Section 4). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Table A.3: Formality and Productivity of AI (Accounting for Organizational Change)

	Support: time		Support: quality		Support: quantity	
	(1)	(4)	(2)	(5)	(3)	(6)
Formal user	0.421*** (0.0775)	0.338*** (0.0663)	0.270*** (0.0829)	0.239*** (0.0739)	0.273*** (0.0776)	0.300*** (0.0654)
Mean of dep. var.	1.948	1.948	2.053	2.053	2.274	2.274
SD of dep. var.	1.270	1.270	1.205	1.205	1.311	1.311
Observations	4,496	4,492	4,436	4,432	4,452	4,448
R-squared	0.225	0.312	0.230	0.318	0.276	0.365
Additional controls	✓	✓	✓	✓	✓	✓
Fixed effects		✓		✓		✓

Notes: This figure shows the results from weighted multivariate regression regressing whether respondents report an improvement in (i) time or speed, (ii) quantity, or (iii) quality of work due to the use of AI on an indicator equaling one if the respondent's main AI application was formally introduced by the employer and zero otherwise. All regression control for a set of individual, establishment, and regional variables. Columns (2), (4), and (6) additionally control for 2-digit industry and 2-digit occupation fixed effects (see Section 4). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1. The sample is restricted to individuals that use AI at their workplace in 2024.

Figure A.3: Formality and Associated Factors after Controlling for Establishment Fixed Effects



Note: This figure shows associations between formal AI introduction and AI usage intensity, AI application type, perceived productivity, AI-based supervision, and training participation, conditional on being an AI user. Estimates are based on weighted linear probability models and control for individual, establishment, and regional characteristics, as well as 2-digit industry, 2-digit occupation fixed effects, and firm fixed effects (see Equation 3 in Section 4). Covariates are measured in 2019 (pre-generative-AI diffusion). 95%-confidence intervals are based on robust standard errors clustered at the establishment level. Observations are weighted using trimmed post-stratification weights described in Section 3.1.

Source: DiWaBe 2.0 Survey



Download ZEW Discussion Papers:

<https://www.zew.de/en/publications/zew-discussion-papers>

or see:

<https://www.ssrn.com/link/ZEW-Ctr-Euro-Econ-Research.html>

<https://ideas.repec.org/s/zbw/zewdip.html>



IMPRINT

**ZEW – Leibniz-Zentrum für Europäische
Wirtschaftsforschung GmbH Mannheim**

ZEW – Leibniz Centre for European
Economic Research

L 7,1 · 68161 Mannheim · Germany

Phone +49 621 1235-01

info@zew.de · zew.de

Discussion Papers are intended to make results of ZEW
research promptly available to other economists in order
to encourage discussion and suggestions for revisions.
The authors are solely responsible for the contents which
do not necessarily represent the opinion of the ZEW.