IHS Working Paper 61

November 2025

AI in Demand: How Expertise Shapes its (Early) Impact on Workers

Eduard Storm

Myrielle Gonschor

Marc Justin Schmidt

INSTITUTE FOR ADVANCED STUDIES VIENNA

Author(s)

Eduard Storm, Myrielle Gonschor, Marc Justin Schmidt

Editor(s)

Robert M. Kunst

Title

AI in Demand: How Expertise Shapes its (Early) Impact on Workers

Funder(s)

DFG

Institut für Höhere Studien - Institute for Advanced Studies (IHS)
Josefstädter Straße 39, A-1080 Wien
T +43 1 59991-0
www.ihs.ac.at
ZVR: 066207973

AI in Demand: How Expertise Shapes its (Early) Impact on Workers*

Eduard Storm[1,2] Myrielle Gonschor[3] Marc Justin Schmidt[4]

November 2025

Abstract

We study how artificial intelligence (AI) affects workers' earnings and employment stability, combining German job vacancy data with administrative records from 2017–2023. Identification comes from changes in workers' exposure to local AI skill demand over time, instrumented with national demand trends. We find no meaningful displacement or productivity effects on average, but notable skill heterogeneity: expert workers with deep domain knowledge gain while non-experts often lose, with returns shaped by occupational task structures. We also document AI-driven reinstatement effects toward analytic and interactive tasks that raise earnings. Overall, our results imply distributional concerns but also job-augmenting potential of early AI technologies.

Keywords: AI, Online Job Vacancies, Skill Demand, Worker-level Analysis, Employment, Earnings, Expertise.
JEL Codes: D22, J23, J24, J31, O33.

[1]Institute for Advanced Studies (IHS). Email: eduard.storm@ihs.ac.at.
[2]RWI – Leibniz Institute for Economic Research.
[3]Kienbaum Consultants.
[4]TU Dortmund, RTG 2484. Email: marcjustin.schmidt@tu-dortmund.de.
**Corresponding author*: Eduard Storm, Institute for Advanced Studies Vienna (IHS); email: eduard.storm@ihs.ac.at. We thank Palturai GmbH / Finbot AG for providing the raw job vacancy data, Niklas Benner for valuable contributions to its preparation for economic analysis as well as Heike Ermert and Alina Niemann for excellent research assistance. We also thank Ronald Bachmann, Karim Bekhtiar, Michael J. Böhm, Lukas Buchheim, Matias Cortes, Gökay Demir, Christina Gathmann, Albrecht Glitz, Terry Gregory, Ines Helm, Andreas Lichter, Fabien Petit, Martin Popp, Duncan Roth, Michael Stops, Simon Wiederhold, Erwin Winkler, and participants from various seminars and conferences for valuable comments and thoughtful discussions from which this paper has benefited greatly. All remaining errors are our own. We gratefully acknowledge financial support from the German Research Foundation (DFG) under the SPP 2267 scheme "Digitalisation of Working Worlds", grant number 670531.

1 Introduction

A growing number of firms and workers are becoming exposed to AI technologies. Unlike previous technologies, AI reshapes domain-specific knowledge —*expertise* —by broadening access to high-skilled positions and enabling more workers to perform tasks that previously required specialized knowledge (Autor 2024; Autor and Thompson 2025). In fact, AI has already affected the task structure of jobs, resulting in establishment-level employment growth for high-skilled jobs, but no aggregate employment effects yet.[1] Micro-level evidence on individual labor market outcomes, however, remains scarce and leaves distributional implications underexplored. Existing worker-level studies use patent- or task-based measures of potential AI exposure, which capture where AI could be adopted rather than where it already shapes labor demand.[2] From a policy-perspective, however, *realized* exposure is crucial to identify who benefits from the emergence of AI and who faces displacement risks.

To address this gap, we examine how workers' earnings and employment outcomes respond to realized changes in AI skill demand. We identify AI skills from German online job vacancies (OJV) and combine the resulting indicator with administrative records. Germany offers a suitable setting for our research question for two reasons. First, it ranks among the world's most innovative economies, with the share of AI-adopting firms rising from 6% in 2019 to 13% in 2023 (Falck, Kerkhof, and Wölfl 2024; Rammer 2022) and 45%—62% of workers reporting AI use.[3] Second, Germany offers high-quality labor market data with detailed information on workers' earnings, employment, and occupational histories, allowing us to precisely estimate AI-induced implications and their underlying drivers.

We study worker-level employment and earnings responses associated with changes in

[1] See Acemoglu, Autor, Hazell, and Restrepo (2022) and Peede and Stops (2024).

[2] See Fossen and Sorgner (2022), Gathmann, Grimm, and Winkler (2024), and Ozgul, Fregin, Stops, Janssen, and Levels (2024).

[3] These figures refer to pre-GenAI technologies. Around 8% of firms in the EU-27 used (pre-Gen) AI technologies in 2023 (Eurostat 2024), with comparable adoption rates in the US (Bonney et al. 2024). Survey evidence further suggests 45% – 62% of German workers use AI tools (Arntz et al. 2024a; Giering et al. 2021). In both Germany and the US, worker-level exposure exceeds firm-level adoption, as large firms account for most AI diffusion (Acemoglu et al. 2025; McElheran et al. 2024).

AI skill demand in their respective local labor market (LLM) —defined as the combination of 3-digit occupations and commuting zones (CZ). Taking advantage of our granular data, we merge this detailed AI exposure measure with administrative worker-level data at the LLM-level. Identification relies on within-LLM changes in AI exposure over time, controlling for individual unobserved heterogeneity, year-specific shocks, regional and occupational productivity differences, and a rich set of worker-and job-related characteristics. Recognizing endogeneity concerns due to non-random AI exposure, we construct a leave-one-out-mean (LOOM) instrument that excludes demand for a worker's own occupation in her CZ, thereby exploiting variation in national AI demand that is plausibly orthogonal to local conditions (Azar, Marinescu, and Steinbaum 2022). Based on balancing checks and placebo tests (Borusyak, Hull, and Jaravel 2025), our instrument is valid as long as we account for occupation-specific demand shocks via 2-digit occupation-by-year fixed effects (FE).

Our analysis is guided by three predictions that are derived from the Acemoglu, Autor, Hazell, and Restrepo (2022) model — henceforth AAHR. This model is characterized by a key dichotomy: AI technologies spur further automation of tasks with negative consequences for the affected workers (*displacement effect*), but they also permit a more flexible task allocation and thus enable workers to become more productive (*productivity effect*). Beyond this dichotomy, AI may also generate entirely new tasks (*reinstatement effect*), creating additional complementarities with human labor. The relative strength of these forces depends on workers' exposure to AI, which varies with their skill endowment and their jobs' task structure. We empirically test these channels and present three main findings.

In our first prediction, we argue that workers facing rising AI exposure in their LLM experience lower employment stability, measured as annual calendar days employed, and limited earnings gains. Our preferred IV specification with 2-digit occupation-by-year fixed effects, however, shows no robust relationship between AI skill demand and workers' outcomes. Viewed through the lens of the AAHR framework, this finding suggests no economically meaningful productivity or displacement effects for workers on average.

Second, we argue that AI primarily complements complex, knowledge-intensive work (Autor 2024). Expert workers —those with deep, domain-specific knowledge —should thus experience higher earnings gains from AI exposure, with weaker implications for employment stability (Autor and Thompson 2025). For this analysis, we interact AI skill demand with codified job complexity measures, thereby distinguishing between four skill groups: helpers, professionals, specialists, and experts. Indeed, we find supporting evidence for our second prediction: a doubling in the share of AI vacancies implies a moderate earnings increase among expert workers by 0.65%. Evaluated at their average real annual earnings, this increase translates to an AI-induced gain of 403 EUR per year. In contrast, non-experts face earnings declines of up to -0.3%. Likewise, we show that the association between AI skill demand and employment stability is increasing in the underlying skill level.

Overall, expert workers consistently realize the largest AI-induced earnings gains, with the size of these gains shaped by the occupational task structure. While experts in cognitively demanding occupations tend to benefit the most from AI skill demand, we find suggestive evidence that AI partly substitutes for analytic tasks in occupations where they are most intensively used (e.g., software developers). Instead, expert gains are more pronounced in cognitive-intensive occupations with a more balanced task composition (e.g., occupations in accounting, controlling, and auditing) and interactive occupations (e.g., sales).

Our third prediction posits that AI creates new tasks, thereby enabling complementarities with labor and amplifying positive outcomes for workers exposed to task space expansion. To quantify this *reinstatement effect*, we follow Deming and Noray (2020) and AAHR by constructing a *net* skill change measure that captures the difference between emerging and fading (non-AI) skills in OJV data. This proxy for changes in workers' relevant task space reveals that AI exposure is associated with an expansion of analytic and interactive activities, but a contraction in manual ones. The reinstatement effect for these cognitively demanding activities generates earnings premia of up to 0.8%, complementing the productivity-enhancing effects of AI —primarily for workers in cognitive-intensive occupations.

Taken together, our analysis shows high-skilled workers benefit in terms of employment stability and earnings, while lesser-skilled workers face higher displacement risks. These findings cast doubt on optimistic views of AI as a potential leveler for reducing inequality (Autor 2024). To assess distributional implications directly, we estimate group-specific responses by assigning workers to deciles of the earnings distribution. For the lowest decile, a doubling in the share of AI vacancies reduces annual working days by eight and earnings by 3.9%, whereas workers in higher deciles gain up to five working days and 2.5% higher earnings. While suggestive, these findings echo concerns about distributional consequences of AI (Acemoglu 2025) and highlight its potential to deepen existing labor market inequalities.

Our paper makes four contributions to the literature on the labor market implications of AI. First, our worker-level analysis adds a micro perspective on earnings and employment responses, complementing studies with industry-, occupation-, and region-level outcomes.[4] Recent studies have shifted attention to firm- and establishment-level responses[5], but worker-level evidence remains rare. Existing studies rely on survey data (Fossen and Sorgner 2022), broader digital technology measures (Genz, Gregory, Janser, Lehmer, and Matthes 2021) or patent-based measures of *potential* AI exposure (Gathmann, Grimm, and Winkler 2024; Ozgul, Fregin, Stops, Janssen, and Levels 2024). Using vacancy-based indicators of *realized* AI demand linked to administrative records, we uncover that modest average responses mask distinct heterogeneity across workers with different skill profiles.

Second, we identify the sources of this skill heterogeneity by linking AI exposure to measures capturing workers' job complexity. Building on Autor (2024) and Autor and Thompson (2025), we show that these job-specific expertise measures provide sharper insights into skill

[4]See Albanesi, Dias da Silva, Jimeno, Lamo, and Wabitsch (2024), Alekseeva, Azar, Giné, Samila, and Taska (2021), Arntz, Genz, Gregory, Lehmer, and Zierahn-Weilage (2024b), Aum and Shin (2025), Bonfiglioli, Crinò, Gancia, and Papadakis (2025), Gathmann and Grimm (2022), Marguerit (2025), Prytkova, Petit, Li, Chaturvedi, and Ciarli (2024), and Webb (2020).

[5]See Acemoglu, Autor, Hazell, and Restrepo (2022), Babina, Fedyk, He, and Hodson (2024), Babina, Fedyk, He, and Hodson (2023), and Hampole, Papanikolaou, Schmidt, and Seegmiller (2025) for the US, Stapleton, Copestake, and Pople (2025) for India, Peede and Stops (2024) for Germany, Aghion, Bunel, Jaravel, Mikaelsen, Roulet, and Søgaard (2025) for France, and Engberg, Hellsten, Javed, Lodefalk, Sabolová, Schroeder, and Tang (2025) for Sweden.

heterogeneity than standard education or human capital proxies. This approach captures meaningful within-occupation variation in skill requirements and helps explain why many studies, including our own baseline, find no discernible implications of AI: gains are concentrated among expert workers, who represent only 16% of the full-time workforce. Accounting for this heterogeneity also helps reconcile mixed findings in earlier worker-level studies, some of which report gains mainly for college-educated workers (Fossen and Sorgner 2022; Gathmann, Grimm, and Winkler 2024), while others find little variation by skill (Ozgul, Fregin, Stops, Janssen, and Levels 2024). By showing that the payoff to skill groups is shaped by their jobs' task structure, we also underscore the need to consider both dimensions when identifying "winners" and "losers" of AI.

Third, we provide worker-level evidence on how AI reshapes the task space through the *reinstatement effect*. Using established skill turnover measures, we quantify the creation of new tasks in online job ads and show that expansions in analytic and interactive activities amplify earnings gains for workers exposed to such task shifts. These insights highlight a job-augmenting mechanism central to modern workhorse models (Acemoglu and Restrepo 2019; Autor, Chin, Salomons, and Seegmiller 2024), but with little empirical micro-level evidence, except notable exceptions (Gathmann, Grimm, and Winkler 2024). While much of the AI literature focuses on automation, our findings also highlight its potential as an augmenting technology that can raise the value of human expertise through task space expansions.

Finally, we contribute new time-varying indicators of AI demand at the occupation and commuting-zone level. Existing indicators are widely used but largely static and US-based,[6] which limits their generalizability (Lewandowski, Madoń, and Park 2025). Our soon-to-be-released indicators, based on the German occupational classification, can be mapped to ISCO and other international systems, enabling harmonized cross-country comparisons.

The paper proceeds as follows. Section 2 outlines our conceptual framework. Section 3

[6] See Brynjolfsson, Mitchell, and Rock (2018), Felten, Raj, and Seamans (2021) and Felten, Raj, and Seamans (2023), and Webb (2020) and, for Europe, Tolan, Pesole, Martínez-Plumed, Fernández-Macías, Hernández-Orallo, and Gómez (2021) and Engberg, Görg, Lodefalk, Javed, Längkvist, Monteiro, Nordås, Schroeder, and Tang (2024).

presents the data, while section 4 outlines details on our AI measure. Section 5 describes the empirical methodology, whose main results and robustness tests are reported in section 6. Section 7 explores skill heterogeneity and reinstatement effects as part of the mechanism analysis, while section 8 discusses inequality implications. Finally, section 9 concludes.

2 Conceptual Framework and Empirical Predictions

This section sketches the key insights from the Acemoglu, Autor, Hazell, and Restrepo (2022) model that guides our empirical analysis. We extend their establishment-level perspective to the worker level. In our setting, establishments' task reallocation decisions are reflected in changes in local AI skill demand. We use these changes as a proxy for shifts in labor demand that affect individual employment stability and earnings (Acemoglu and Restrepo 2022).

The AAHR Model: Displacement vs. Productivity Effect. AAHR assume a competitive market where profit-maximizing establishments combine labor and AI technologies to produce output. AI is modeled as a productivity-enhancing technology that can substitute for labor in existing tasks. As AI becomes more effective, firms expand the set of tasks assigned to AI. This *displacement effect* reduces labor demand. Yet, automation also lowers costs and thus allows for more flexible task allocation. This *productivity effect* raises labor demand. The net effect on labor demand depends on the relative size of these two forces and varies with the share of tasks that can be profitably automated.

Empirical Expectations at the Worker-Level. First, we examine the implications of AI exposure for the average worker. In the AAHR model, the displacement effect dominates unless productivity gains are sufficiently large. Establishment-level evidence, however, points to modest productivity effects and no significant net job creation to date (Peede and Stops 2024). We thus expect displacement to dominate for employment outcomes. In contrast, implications for earnings are less clear. While displacement reduces labor demand, productivity gains can raise the value-added of remaining jobs. We thus expect average implications on earnings to be muted or small at best. This reasoning leads to our first prediction:

Prediction 1: Displacement Effect dominates over Productivity Effect

Workers exposed to AI technologies experience lower employment stability and (at most) limited earnings gains on average.

We test Prediction 1 by regressing worker-level outcomes on the share of vacancies in their local labor market that require AI skills. Positive coefficients indicate that productivity effects dominate, while negative coefficients indicate that displacement dominates.

Second, AAHR assume perfect substitutability between AI and labor, thus abstracting from worker heterogeneity. Yet, AI differs from earlier technologies in that it may extend, rather than replace, workers' domain-specific knowledge (Autor 2024). This *expertise* shapes how workers are affected: expert workers are better positioned to enhance their productivity by leveraging AI tools, while non-expert workers face greater displacement risks. Autor and Thompson (2025) formalize this mechanism, showing that automation tends to raise wages but lower employment when it increases expertise requirements (and vice versa when expertise requirements decline). Following this logic, we argue in our second prediction that workers with more expertise are the primary beneficiaries of AI:

Prediction 2: AI complements Worker Expertise

Workers in more complex jobs experience higher earnings gains from AI exposure, with comparably weaker implications on employment stability.

We test Prediction 2 by interacting AI exposure with worker skill groups. To operationalize skill groups, we use codified measures of job complexity that distinguish between helpers, professionals, specialists, and experts.

Third, while AAHR emphasize displacement and productivity effects in their theoretical framework, new technologies can create entirely new tasks via the *reinstatement effect* (Acemoglu and Restrepo 2019). Expanding the task space generates new skill requirements that, due to initial scarcity, should expand employment opportunities and yield earnings gains.

Our final prediction thus posits that AI reshapes labor demand also by creating new tasks:

Prediction 3: Reinstatement Effect amplifies worker outcomes

AI-induced creation of new tasks increases employment opportunities and earnings for workers exposed to expanding skill requirements.

We test Prediction 3 by tracking the net skill change in a worker's relevant task space, based on a comparison of emerging versus fading skills in online job ads. Interacting AI exposure with this net skill change measure allows us to assess the extent to which expansions in the task space amplify earnings and employment outcomes.

3 Data: Online Job Vacancies & Administrative Records

In this section, we present our data and the steps to construct our AI exposure measure. We use German online job vacancies from 2017–2023 to identify AI skills and a 2% representative sample of the German working population to construct our outcome and control variables.

Online Job Vacancies. Job postings are collected by our partner —Finbot AG, an IT-company from Meerbusch, Germany. Finbot is a subsidiary of Palturai GmbH, from Hofheim, Germany, and offers custom-made firm-, person- and job posting- data and market analyses. To this end, they scrape vacancies from job boards, company websites, temporary employment agencies, and head-hunters. Finbot consistently updates their online sources and scrapes all sources on a daily basis. Subsequently, Finbot performs basic cleaning procedures and removes duplicates from the same source, before sharing the data with us.

Our OJV data offers two key features compared to other vacancy data commonly used in economic research. First, we have access to the original job vacancies, including all texts included in the posting. This feature access allows us to have more control over data preparation and develop our own, transparent taxonomies. Second, our data access allows us to divide vacancies into segments with distinct purposes.

A typical job vacancy comprises four major sections: i) firm description, introducing the employer, ii) job description, describing the role and responsibilities, iii) skill requirements, and iv) benefits and perks associated with the advertised position. Our access to the raw text data allows us to precisely identify AI skills from the job description and skill requirements segments. In our view, these are the relevant segments for the identification of AI skills. Our data thus ensures that our AI measures are rooted in the specific context of each posting. In Appendix A1.1 we provide further details on our data preparation.

A key concern with vacancy data is non-representativeness as job ads tend to be tilted towards high-skilled jobs. We thus perform weighting exercises to align the data closer with representative employment structures. In Appendix A1.2 we discuss the external validity of our OJV data, comparing it to data on open positions from surveys and official statistics. We show that weighted and unweighted data deliver similar insights on stylized facts on AI skill demand that we exploit for our empirical analysis. The skewness inherent in our OJV data is thus unlikely to distort our empirical analysis.

Eventually, we use 7.7 million vacancies for our analysis. Our final data only comprises vacancies advertising regular work, i.e. full- or part-time, thereby removing vacancies seeking apprenticeships, trainees, and other types of irregular work. To ensure high data quality and maintain comparability with our administrative data, we also remove vacancies (i) posted by temporary employment agencies and (ii) those missing information on key variables.

Administrative Data. We use the Sample of Integrated Labor Market Biographies (SIAB v1723), a 2% representative sample of administrative data on all workers subject to social security contributions (SSC) for the period 1975-2023.[7] As is common in administrative data, wage information is top-censored for around five percent of all spells (Dauth and Eppelsheimer 2020). Following standard procedures, we use imputations for education and wages provided by the IAB-FDZ (Fitzenberger, Osikominu, and Völter 2006). To analyze

[7]The SSC requirement excludes certain individuals, notably the self-employed and civil servants. See Schmucker and Berge (2025) and Graf, Grießemer, Köhler, Lehnert, Moczall, Oertel, Schmucker, Schneider, and Berge (2025) for details on the data.

real annual earnings, we deflate imputed (daily) wages by the national consumer price index and multiply them with the number of days employed. We focus on full-time workers aged 18-65 and exclude workers (i) with zero wage and wages below the first percentile, (ii) from agricultural, forestry and fishery occupations and (iii) those with any missing information on our outcome (employment, earnings) and control variables, including observations from LLMs with too few observations (due to data privacy reasons).

We further supplement the SIAB with the Establishment History Panel (BHP), which covers all establishments in the IAB employment history, collecting information on employment, industry, and work location of establishments between 2012–2023. For parts of our mechanism analysis, we also draw on self-generated non-AI skill measures based on BERUFENET, an online job portal maintained by the German Federal Employment Agency. BERUFENET provides detailed information on job requirements —covering tasks, skills, certificates, and technologies —and is conceptually comparable to O*NET and ESCO.

4 Measurement of AI Exposure

Next, we describe the construction of our AI measure and key stylized facts on differential exposure across workers. We identify AI skills from job descriptions in OJV, combining a keyword-based approach with manual annotation and assistance by ChatGPT 4.0. For brevity, we outline key data preparation steps and refer to Appendix A2.1 for details.

Identification of AI skill demand. In the first step, we compiled an initial list of 97 AI skills based on keywords used in previous studies, including AAHR.[8] In the second step, we manually annotated a random sample of job postings to validate and refine this list, since the initial set lacked references to recent innovations, specific AI tools (e.g., Python packages), and terminology common in German job vacancies. Two student workers independently

[8] See also Büchel, Demary, Goecke, Kohlisch, Koppel, Mertens, Rusche, Scheufen, and Wendt (2021), Hunt, Cockburn, and Bessen (2024), and Taska, O'Kane, and Nania (2022).

classified AI skills following a predefined set of annotation rules to ensure consistency. In cases of disagreement, the author team made the final classification decision. We applied a conservative approach, generally excluding borderline cases from the final list.

In total, 275 AI skills enter our final taxonomy, extending previous taxonomies by incorporating newly emerging AI skills and distinguishing between distinct domains.[9] We differentiate between two broad domains: (i) *AI methods*, reflecting developer-stage tools such as algorithms, model families, and software libraries, and (ii) *AI applications*, which capture industry-specific uses such as autonomous driving or image/ voice recognition. As expected, *AI methods* appear more frequently in job postings, as *AI applications* have a narrower scope (Appendix A2.2, Figure B3).[10] Note that demand for GenAI skills is negligible to date, thus not meaningfully captured in our measure. While we include AI skills related to GenAI, our measure must be interpreted as capturing "pre-GenAI" skills.

We validate our AI skill measure against widely-used definitions from the literature, including other vacancy-based sources (e.g., Lightcast, Indeed) and the Felten, Raj, and Seamans (2021) measure. We discuss this external validity in detail in Appendix A2.3.

Finally, we measure AI exposure, AI_{lot}, by computing the share of vacancies requiring AI skills in a worker's LLM, i.e. $AI_{lot} = \frac{N_{lot}^{AI}}{N_{lot}}$. The numerator, N_{lot}^{AI}, captures the number of vacancies listing at least one AI skill, while the denominator, N_{lot}, includes all posted job ads in a worker's LLM. We define AI exposure at the LLM-level to accommodate a merge with administrative data. Thanks to our granular data, we adopt a detailed LLM definition at the CZ-occupation-year level. At the regional level, we aggregate 402 counties into 141 CZs (Kosfeld and Werner 2012). At the occupational level, we exclude postings in the armed forces, agriculture, fishing, and forestry, and retain only LLMs with at least three postings

[9] In Appendix A2.1 we illustrate all AI skills (Figure B2) and list the Top 40 AI skills (Table B1).

[10] This tilt toward *AI methods* is to be expected at this early stage of AI adoption. However, our baseline measure still provides meaningful insights into AI-induced changes in labor demand, which is also reinforced by firm-survey evidence. For example, Rammer, Fernández, and Czarnitzki (2022) document the widespread use of machine learning and text mining in product development. Even when firms rely on external AI systems, effective use requires expertise in data handling, API integration, and model oversight, which should be plausibly reflected in job postings (Falck, Kerkhof, and Wölfl 2024).

Figure 1: Trends in AI skill Demand in Germany, 2017–2023

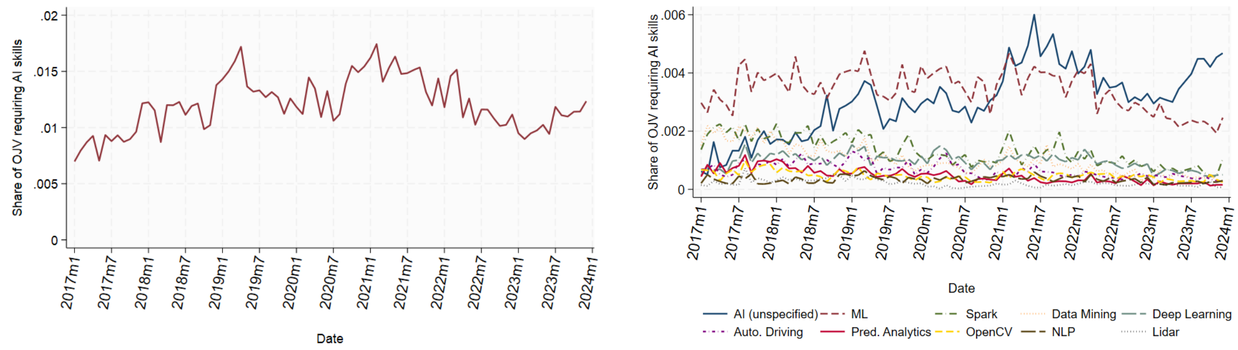

(a) Aggregate AI skill demand

(b) Top 10 AI Skills

NOTE. — Panel 1a depicts the aggregate share of AI vacancies, while Panel 1b shows the Top 10 skills. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; own calculations.

per year. These restrictions leave us with 131 occupations and 141 CZs, implying 18,471 LLMs for which we compute year-specific AI exposure measures. We apply a June 30 cutoff each year to align with the cutoff date in our administrative data.

AI Skill Demand over Time and by Skill Groups. Figure 1a displays the share of AI vacancies from 2017 to 2023. Despite notable short-term fluctuations, aggregate AI skill demand has shown no sustained upward trend: the share of AI vacancies increased from around 1% in 2017 to 1.5% in 2021, before returning to 1% by 2023. Panel 1b shows most of AI skill demand reflects demand for unspecified AI skills and Machine Learning, rather than more specialized domains such as deep learning or NLP. Our enriched AI skill measure captures similar dynamics to other measures from the literature, but with more stable trends over time and deeper insights into AI-domain-specific demand shifts (Aghion, Bunel, Jaravel, Mikaelsen, Roulet, and Søgaard 2025), as discussed in Appendix A2.3.

While aggregate AI skill demand has remained relatively stable over time, the cross-sectional distribution of AI exposure across workers shows substantial heterogeneity. Table 1 documents a clear skill gradient: high-skilled workers face substantially higher AI exposure. AI exposure is highest among workers in analytic task-intensive occupations (2.2%), expert

Table 1: AI Skill Demand, Days Employed, and Earnings by Skill Groups

	Share in Sample	AI OJV Share (%)	Days Employed (Yr.)	Earnings (EUR)
No degree	0.07	0.34	345.2	32,234
Vocational	0.69	0.37	357.0	41,232
College	0.24	1.34	356.7	58,805
Helpers	0.11	0.10	347.5	30,153
Professionals	0.55	0.34	356.7	39,918
Specialists	0.17	0.96	358.3	53,357
Experts	0.17	1.45	357.7	61,967
NR manual	0.15	0.05	352.7	33,928
Routine manual	0.15	0.11	353.3	35,913
Routine cognitive	0.42	0.32	357.7	45,774
NR interactive	0.09	0.31	356.5	48,778
NR analytic	0.19	2.21	357.5	57,274
Full sample	1.00	0.61	355.0	44,908
Observations	2,455,816	2,455,816	2,445,816	2,455,816

Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB v7523 (Earnings and employment data); IAB Occupational Panel (Task group data); own calculations.

workers (1.4%), and college graduates (1.3%). Notably, these skill groups also work more days per calender and earn higher average annual earnings, suggesting that labor market implications of AI should be most pronounced among these high-skilled workers.

5 Methodology

Baseline Specification. We now test the implications of AI exposure on worker' outcomes. Our first prediction states that workers exposed to AI face lower employment stability and limited earnings gains (section 2). We first consider a simple model using OLS:

$$Y_{ilot} = \alpha_i + \beta_{1,OLS} AI_{lot} + \beta_2 X_{it} + \psi_l + \omega_o + \theta_t + \epsilon_{ilot} \tag{1}$$

Here, Y_{ilot} denotes our outcome variables for worker i, working in CZ l, occupation o, and year t: (i) L_{ilot}, number of calendar days employed in t; and (ii) E_{ilot}, log annual earnings. Our key regressor AI_{lot} measures AI exposure as the share of vacancies requiring AI skills in the worker's LLM. We control for a rich set of worker-level covariates X_{it} (age, age squared, education, gender, skill level, foreign origin, firm tenure, firm size, 2-digit industry (WZ08)). Furthermore, we include worker FE (α_i) to capture unobserved heterogeneity, year FE (θ_t)

Figure 2: Variation in AI Skill Demand: Occupation-CZ Level, 2017–2023

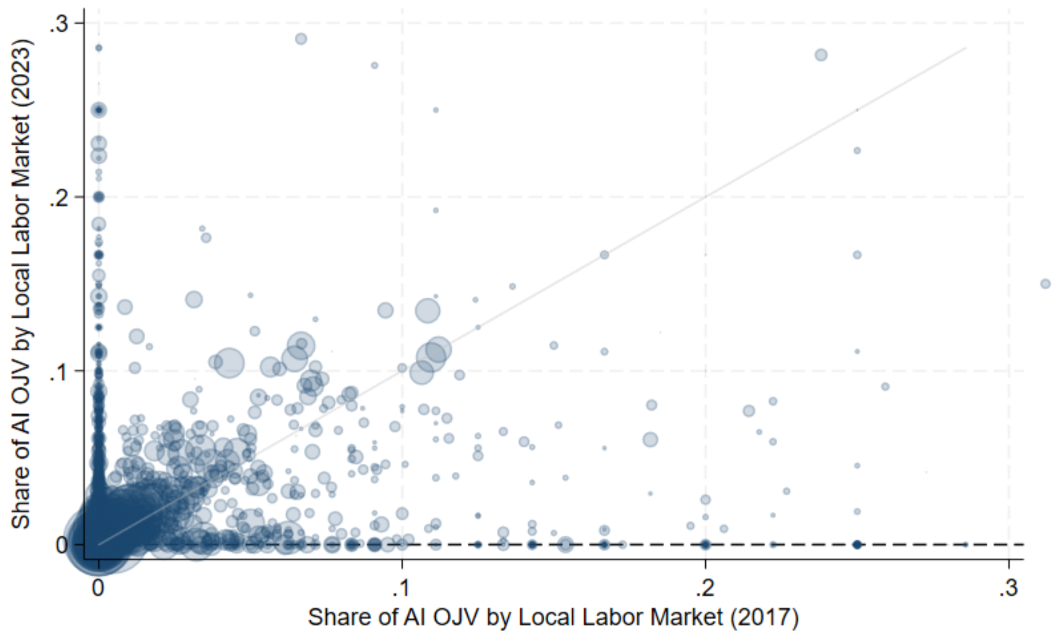

NOTE. —The X-axis shows the share of online job vacancies (OJVs) requiring AI skills in 2017, while the Y-axis shows the corresponding shares in 2023. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; own calculations.

for common shocks, and separate CZ FE (ψ_l) and 3-digit occupation FE (ω_o) to account for systematic differences in productivity and technology adoption.

Our research design compares workers within the same occupation and region while allowing AI_{lot} to vary at the occupation–CZ–year level. Identification thus requires employment and earnings outcomes to follow common time, occupation, and regional trends.

Figure 2 highlights this variation, plotting the share of AI vacancies across LLMs in 2017 against 2023. AI exposure varies widely across LLMs and indicates that early AI exposure shapes longer-term dynamics. The identifying variation in our research design comes from changes in AI demand in a worker's LLM over time. While most LLMs display little AI skill demand throughout, others experienced notable increases. For example, in 2017 only 9% of LLMs displayed *some* kind AI skill demand. By 2023, this share nearly doubled to 16%, reflecting meaningful variation in AI exposure over time.

Identification Threats and IV Strategy. Despite efforts to control for confounding factors, OLS may suffer from endogeneity. On the one hand, reverse causality may arise if higher earnings attract workers with suitable skills, thus raising (local) labor supply. Holding other factors constant, this channel puts downward pressure on earnings.[11] On the other hand, higher earnings reduce labor demand via the job creation curve (Bassier, Manning, and Petrongolo 2025; Lichter, Peichl, and Siegloch 2015), which mechanically reduces the number of potential jobs with AI skill demand. Because of these downward pressures, both channels imply that the OLS coefficient $\widehat{\beta}_{1,OLS}$ is likely biased downward.

To address these identification threats, we exploit occupation-level shifts in national AI skill demand with a leave-one-out-mean (LOOM) instrument:

$$AI_{lot}^{IV} = \frac{1}{L-1} \sum_{l' \neq l}^{L} AI_{l'ot} \tag{2}$$

By excluding occupation-specific demand in worker's own CZ l, this instrument removes the direct influence of local shocks, thereby isolating variation in national changes in AI demand within occupation o and year t that is plausibly orthogonal to local conditions (Azar, Marinescu, and Steinbaum 2022). Considering we distinguish between 131 occupations, we draw on a relatively large number of shocks, in line with quasi-experimental conditions for consistency (Borusyak, Hull, and Jaravel 2022; Borusyak, Hull, and Jaravel 2025). We interpret these shifts as technology shocks that are plausibly exogenous from the perspective of individual workers. For example, a rise in national AI skill demand in a given occupation may reflect innovations trends originating outside of Germany, such as advancements in Machine Learning frameworks (e.g., TensorFlow, PyTorch), computer vision, or the emergence of Large Language Models (Devlin, Chang, Lee, and Toutanova 2019).

Our LOOM design is closely related to shift-share instruments often used in the literature on technological change (Acemoglu and Restrepo 2020; Dauth, Findeisen, Suedekum, and Woessner 2021b) or globalization (Autor, Dorn, and Hanson 2013; Dauth, Findeisen, and

[11] Moreover, if this channel increases the average skill level of workers, firms may be more likely to adopt AI technologies as suggested by the literature on directed technological change (Acemoglu 2002).

Suedekum 2021a). The key distinction is that we have no need for apportioning of national shocks (*shift*) to regions (*share*), as is common in these (region-level) studies. By exploiting occupation-level shocks directly, our LOOM correction ensures that national AI demand for occupation o is purged of mechanical correlation with our outcomes. Our design is thus closely related to studies where worker-level outcomes are regressed on (leave-one-out) national trends, such as trade shocks (Autor, Dorn, Hanson, and Song 2014) or labor market tightness (Börschlein, Bossler, and Popp 2024). We also use a conventional shift-share design for robustness, which provides similar insights (section 6.3).

Balancing Tests and Placebo Analysis. While LOOM instruments usually satisfy the relevance condition —which we confirm with strong first-stage F-statistics —they remain vulnerable to spillovers and predating trends. Following recent econometric recommendations, we thus test the validity of shock orthogonality with balancing checks and placebo tests (Borusyak, Hull, and Jaravel 2022; Borusyak, Hull, and Jaravel 2025). Balancing checks show that our instrument is not systematically correlated with a broad set of predetermined individual, occupational, and firm characteristics (Table B3 in Appendix A3.). We include these controls throughout our main regressions to account for potential residual imbalances.

Our main validation comes from placebo tests. Here, we regress our earnings and employment outcome variables from 2012–2016 on AI skill demand between 2017–2021, such that outcomes in 2012 are regressed on AI demand in 2017, and so on.[12] Under the exclusion restriction, $\widehat{\beta}_1$ in this regression should be zero. In Appendix A3, Table B4, we show that employment regressions yield insignificant estimates, suggesting no meaningful pre-trends. However, earnings regressions show significant coefficients in less saturated models, which vanish once (2-digit) occupation-by-year FE are included. Guided by these results, we incorporate time-varying occupation-specific shocks in our preferred specification.

[12] We restrict the analysis to AI exposure between 2017–2021 for two reasons. First, a structural break in the classification of occupations prevents us from using more distant worker-level information. Second, we want to avoid overlap with the time horizon of our main analysis (2017–2023). In an alternative specification, we regress predetermined worker-level outcomes from 2012–2016 on average AI exposure from 2017–2023, which provides qualitatively similar insights.

Preferred 2SLS model. For our main analysis we estimate a two-stage least squares (2SLS) model, which incorporates the preferred FE structure implied by our placebo analysis:

$$AI_{lot} = \pi_1 AI_{lot}^{IV} + \pi_2 X_{it} + \psi_l + \omega_{o3} + \theta_t + (\gamma_{o2} \times \theta_t) + \nu_{ilot} \tag{3a}$$

$$Y_{ilot} = \alpha_i + \beta_{1,IV} \widehat{AI}_{lot} + \beta_{2,IV} X_{it} + \psi_l + \omega_{o3} + \theta_t + (\gamma_{o2} \times \theta_t) + \epsilon_{ilot} \tag{3b}$$

In eq. (3a), we first predict local AI exposure using our LOOM instrument. In eq. (3b), we then estimate the impact of predicted AI exposure on our outcomes of interest. We are primarily interested in the sign and magnitude of $\widehat{\beta}_{1,IV}$. Through the lens of our conceptual background (Section 2), this coefficient is informative on the relative size of the displacement and productivity effect. We interpret $\widehat{\beta}_{1,IV} > 0$ as consistent with a relatively strong productivity effect and $\widehat{\beta}_{1,IV} < 0$ as consistent with a comparably strong displacement effect.

As argued before, OLS estimates likely suffer from downward bias —due to higher earnings attracting skilled workers or reducing job creation —which is why we expect $\widehat{\beta}_{1,IV} > \widehat{\beta}_{1,OLS}$. To make results comparable across workers and interpretable in economic terms, we normalize our AI exposure measure by its mean such that it permits an interpretation in response to a doubling in the share of AI vacancies.

6 Baseline Results

6.1 Employment: Number of Calendar Days Employed

In Table 2 we report the results for our employment outcome. Reassuringly, the first-stage F-statistics comfortably exceed conventional thresholds, supporting the relevance of our instrument. All specifications include worker, CZ and 3-digit occupation FE, and progressively add further covariates, including socio-demographics (column 2), work history (column 3), employer characteristics (column 4), and (2-digit) occupation-by-year FE (column 5).

OLS estimates are all small, suggesting no meaningful association between AI exposure and employment stability. IV estimates are statistically significant, but also very minor. In

Table 2: IV Regressions of Employment Stability on AI Skill Demand

	Number of Calendar Days Employed				
	(1)	(2)	(3)	(4)	(5)
AI Skill Demand (OLS)	-0.0304**	-0.0198*	-0.0200*	-0.0182*	-0.0042
	(0.0118)	(0.0107)	(0.0108)	(0.0107)	(0.0105)
AI Skill Demand (IV)	-0.2617***	-0.1469***	-0.1485***	-0.1481***	0.0911*
	(0.0358)	(0.0363)	(0.0363)	(0.0363)	(0.0470)
Controls: Socio		✓	✓	✓	✓
Controls: Work			✓	✓	✓
Controls: Firm				✓	✓
Occ. (2-digit) × Year FE					✓
AI Share Mean			0.006		
F-statistic (1st stage)	3604.0	3253.0	3253.0	3253.1	1301.2
R^2 (OLS)	0.4011	0.4038	0.4040	0.4053	0.4059
Observations	2,385,614	2,385,614	2,385,614	2,385,614	2,385,614

NOTE. —All specifications include worker, CZ, and 3-digit occupation FE. Additional controls are added sequentially: socio-economic (gender, education, age, foreign origin, skill level), work (tenure), and firm (size, industry). Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB v7523; own calculations.

the baseline specifications (columns 1–4), we find small negative coefficients ranging from -0.26 to -0.14, implying that a doubling in the share of AI vacancies reduces annual employment by about a quarter of a day. Once we include occupation-by-year FE (column 5), the point estimate turns slightly positive at 0.09. This sign flip implies that broad occupation-specific time trends absorb shocks correlated with AI demand, leaving only within-occupation variation less likely to be confounded. Economically, however, all estimates suggest that AI exposure has no meaningful implications on employment stability for the average worker.

Using the same administrative data, but with AI exposure derived from patent data, Gathmann, Grimm, and Winkler (2024) report modest displacement effects for the average worker. Their approach identifies AI exposure at the industry level and covers a longer horizon (2004–2021), which gives potential displacement effects more time to materialize. Our results are, however, consistent with establishment-level studies using vacancy-based measures, which likewise find no discernible displacement effects (Peede and Stops 2024).

Table 3: IV Baseline Regressions of Log Annual Earnings on AI Skill Demand

	Log Annual Earnings				
	(1)	(2)	(3)	(4)	(5)
AI Skill Demand (OLS)	0.0003**	0.0001	0.0001	0.0002	0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
AI Skill Demand (IV)	0.0082***	0.0071***	0.0071***	0.0071***	0.0001
	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Controls: Socio		✓	✓	✓	✓
Controls: Work			✓	✓	✓
Controls: Firm				✓	✓
Occ. (2-digit) × Year FE					✓
AI Share Mean			0.006		
F-statistic (1st stage)	3258.9	3251.8	3251.8	3251.9	1301.2
R^2 (OLS)	0.8870	0.8925	0.8925	0.8946	0.8952
Observations	2,385,614	2,385,614	2,385,614	2,385,614	2,385,614

NOTE. —All specifications include worker, CZ, and 3-digit occupation FE. Additional controls are added sequentially: socio-economic (gender, education, age, foreign origin, skill level), work (tenure), and firm (size, industry). Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB v7523; own calculations.

6.2 Annual Earnings

Next, we turn to our results on annual earnings (Table 3). As with the employment analysis, first-stage F-statistics comfortably exceed conventional thresholds across specifications. All models include worker, year, CZ and 3-digit occupation FE, with progressively richer sets of controls, analogous to the discussion of our employment results.

OLS estimates are consistently small and mostly insignificant. As expected, IV estimates are larger in magnitude, but remain modest overall. In the baseline specifications without occupation-by-year FE (columns 1–4), we find positive coefficients of around 0.7—0.8%, suggesting a doubling in the share of AI vacancies implies an earnings increase of less than 1% for the average worker. Once we include occupation-by-year FE (column 5), however, the coefficients become statistically insignificant. This pattern mirrors our employment results in that broad occupation-specific time trends absorb shocks correlated with AI demand, leaving less variation for identification. Combined, the evidence points to, at best, limited earnings gains from AI exposure that vanish entirely once more restrictive controls are included.

6.3 Robustness

We perform various robustness tests to validate our main results. Specifically, we (i) use alternative AI measures, (ii) explore the role of mobility, (iii) test model sensitivity, (iv) account for wage censoring and sample selection, and (v) construct alternative instrument specifications. All tests are based on our main IV specification in eq. (3), but exclude (2-digit) occupation-by-year FE to keep sufficient identifying variation for meaningful comparisons. For brevity, we summarize the main insights here and provide a detailed discussion of our robustness results in Appendix A4.

First, we show that our employment results are similarly driven by developer-stage skills ("AI methods") and application-related skills ("AI applications"), while earnings gains are more pronounced for AI applications (Appendix A4.1, Table C1). Combined, these results suggest that AI domains affect labor demand through distinct channels. In addition, we construct AI measures building upon existing taxonomies from the literature, including AAHR (Appendix A4.1, Table C2). These measures provide similar results, though our AI measure implies the largest earnings gains.

Second, we explore the role of mobility by sequentially relaxing FE that constrain worker movement across CZs or occupations (Appendix A4.1, Table C3). Employment estimates are slightly negative when identified within occupations, but vanish once we allow for worker movement across occupations. Earnings estimates, in turn, remain positive but attenuate once we permit occupational mobility. These patterns suggest job mobility can dampen displacement risks, while earnings gains are concentrated in within-occupation specifications.

Third, we assess sensitivity to alternative scaling choices of our AI measure and non-linear specifications (Appendix A4.1, Table C4). While scaling choices leave our results intact, we find a convex relationship between AI exposure and employment, implying any adverse effects are concentrated at low levels of AI exposure. In contrast, we find a concave relationship between AI exposure and earnings, indicating diminishing gains at very high levels of AI exposure. AI exposure may thus initially benefit workers through higher pay, but only translates into employment gains once demand reaches a sufficiently high level.

Fourth, we account for sample selection (Appendix A4.1, Table C5) and potential biases from wage censoring (Appendix A4.1, Table C6). Results are robust to alternative sample restrictions and trimming of extreme values. Yet, replacing imputed wages with top-coded values turns the earnings coefficient negative, suggesting AI-induced earnings gains primarily benefit high-skilled workers and may thus reinforce income inequality. As wages of high earners would otherwise be censored, this exercise underscores the importance of addressing wage censoring when estimating technology-induced distributional effects.

Finally, we construct a spillover-resistant instrument in which we exclude adjacent CZs (Appendix A4.1, Table C7) and a conventional shift-share IV (Appendix A4.1, Table C8). Results are similar to those obtained with our main instrument, reinforcing that our findings are not merely an artifact of the specific IV construction.

7 Mechanisms

In this section, we examine two mechanisms grounded in our conceptual framework (Section 2). First, we show that higher-skilled workers with deep domain expertise exhibit more positive labor market outcomes due to complementarities with AI (Prediction 2). Second, we demonstrate workers not only benefit from AI through productivity gains associated with existing tasks, but also through the creation of new tasks (Prediction 3).

For these analyses, we focus on our LOOM IV approach and interact the first-stage estimate $\widehat{AI}_{lot}$ with group dummies D_{it}^s, resulting in S instruments for S endogenous variables:[13]

$$Y_{ilot} = \alpha_i + \sum_{s=1}^{S} \beta_{1,s} D_i^s \times \widehat{AI}_{lot} + \beta_2 X_{it} + \psi_l + \omega_{o3} + \theta_t + (\gamma_{o2} \times \theta_t) + \epsilon_{ilot} \qquad (4)$$

where $D_i^s = 1$ if the underlying categorical variable has some value s and $D_i^s = 0$ otherwise. Dummies D_{it}^s and other base effects are absorbed by CZ, 3-digit occupation

[13] We follow Börschlein, Bossler, and Popp (2024) who study heterogeneous effects on the relationship between (worker-level) wages and LLM-specific tightness measures, thus a methodology very similar to ours.

FE, and (2-digit) occupation-by-year FE.[14] As in our baseline IV, identification comes from national shocks to AI demand at the occupation level. The interaction terms $D_i^s \times \widehat{AI}_{lot}$ capture group-specific responses, allowing us to examine differential responses to changes in AI skill demand for distinct skill groups. Unlike traditional approaches, in which the regressor is interacted with group indicators, our approach directly estimates the employment and earnings implications of AI exposure for each group s, without reference to a base group.

For ease of interpretation, we normalize AI exposure by dividing each worker's exposure by the group-specific mean, ensuring that a one-unit increase reflects a doubling in the share of AI vacancies relative to the group mean. This normalization abstracts from level differences across skill groups that would otherwise distort comparisons.[15]

7.1 Skill Heterogeneity

We begin by examining if job-specific expertise determines which workers benefit from AI skill demand. To do so, we draw on codified measures of job complexity in the German occupational classification (KldB 2010 v2020), which distinguishes four groups: helpers, professionals, specialists, and experts. AI exposure differs systematically across these groups, ranging from 0.1% for helpers, 0.3% for professionals, 1% for specialists, and 1.4% for experts. These categories allow us to trace differential AI exposure across rising job requirements within the same occupation. Recognizing that the payoff to expertise may vary with the bundles of tasks workers commonly perform, we also account for occupational task structures.

Job-specific expertise. Figure 3 shows our main results on skill heterogeneity, separately for employment stability and annual earnings. Panel 3a indicates that the association between demand for AI skills and number of days employed rises with the skill level. A dou-

[14]Unlike the robustness analysis, where we relaxed occupation-by-year FE to preserve variation and ensure comparability across specifications, we include them for our mechanism analysis to isolate group-specific responses with more precision, following the implications of our placebo tests. We do, however, find qualitatively similar results when excluding occupation-by-year FE or, alternatively, running OLS regressions.

[15]For example, using a simple share-based measure would inflate coefficients of lesser-skilled workers, since a 1 pp. increase represents a much larger relative change for groups with low baseline exposure.

Figure 3: IV Regressions of Earnings and Employment on AI Skill Demand: By Worker Skill Groups (Job Complexity)

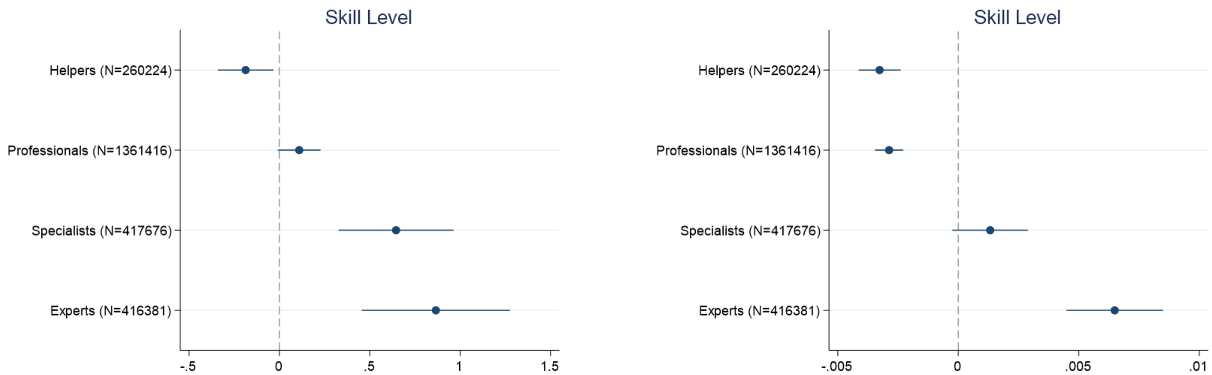

(a) Number of Calendar Days Employed

(b) Log Annual Earnings

NOTE. — All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation FE, and 2-digit-occupation-by-year FE. Earnings regressions also include year FE. Numerical results are reported in Appendix A6, Table E1. Point estimates are shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

bling in group-specific share of AI vacancies increases employment by 0.9 days for experts and specialists (0.25% relative to baseline number of days employed), while lesser-skilled workers lose a quarter of a day (-0.10%). Panel 3b highlights that only experts gain significantly in earnings. A doubling in the share of AI vacancies raises their earnings by 0.65%, which, evaluated at their average annual earnings of 61,966 EUR, translates to an AI-induced net gain of 403 EUR. In contrast, lesser-skilled workers face minor earnings declines of around -0.3%. Overall, we find expertise requirements imply larger earnings gains than conventional skill measures such as formal education (Appendix A6, Table E2) or accumulated human capital via labor market experience (Appendix A6, Table E4).[16]

Supplementary analyses provide further insights into the heterogeneous implications of distinct AI domains. In Appendix A6 (Table E1) we show that employment gains for experts

[16]Mirroring Gathmann, Grimm, and Winkler (2024), we find that only college-educated workers benefit from rising AI demand. College graduates gain both in employment stability (+0.8 days) and earnings (+0.4%), while non-graduates face losses in employment stability (-0.4 days) and earnings (-0.9%). For human capital, we find that junior workers with less than five years of experience exhibit the largest employment gains (+1 day), while experienced workers with up to 15 years of experience see the largest earnings gains (+0.4%). Using either alternative measure, however, yields smaller earnings gains than job complexity.

are driven by "AI Methods". In contrast, expert workers' earnings gains are primarily driven by "AI Applications" such as LIDAR (1.75%), whereas earnings gains associated with AI Methods such as Machine Learning are more modest in comparison. These patterns suggest AI-induced labor demand shifts affect workers differently depending on the domain in which AI is being applied (Aghion, Bunel, Jaravel, Mikaelsen, Roulet, and Søgaard 2025).

Overall, our results support the second prediction: earnings gains are concentrated among expert workers, while employment stability is smaller. These patterns explain the insignificant average estimates in our preferred specification, as positive estimates are limited to expert workers —who represent only 16% of the full-time workforce.

Occupational Task Structure. Having established that expertise pays off for workers exposed to AI, we now examine whether outcomes differ by their occupational task structure. Since workers apply their knowledge to specific bundles of tasks, the payoff to AI exposure may depend on the types of activities they commonly perform.

To capture task structures, we use skill requirements from BERUFENET —the German equivalent of ONET (Dengler, Matthes, and Paulus 2014). We process the raw information into a structured dataset with 9,500 keywords, which we map to non-AI related requirements in our OJV data (see Appendix A5 for details). We then classify these requirements into five established categories: (i) non-routine (NR) analytic, (ii) NR interactive, (iii) routine (R) cognitive, (iv) R manual, and (v) NR manual. For each posting, we compute the relative frequency of skills in each category and aggregate them at the 3-digit occupation level. The resulting shares are proxies for the occupational task space, and we assign each occupation to the task group with the highest share.[17] Because task space is measured at the occupation level, we run a simplified version of eq. (4), including 3-digit occupation FE but excluding 2-digit occupation-by-year FE, which would absorb most of the variation.

Figure 4 shows that earnings responses differ sharply by occupational routine intensity. We find meaningful earnings gains only in analytic (+0.6%) and interactive (+0.4%) occu-

[17] Table D2 in Appendix A5 displays the Top 5 occupations by task intensity to illustrate this classification.

Figure 4: IV Regressions of Earnings and Employment on AI Skill Demand: By Occupational Task Structure

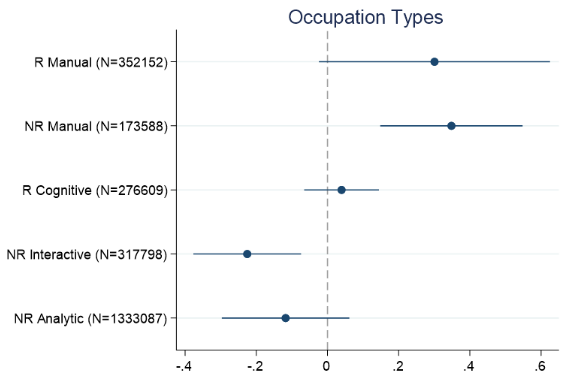

(a) Number of Calendar Days Employed

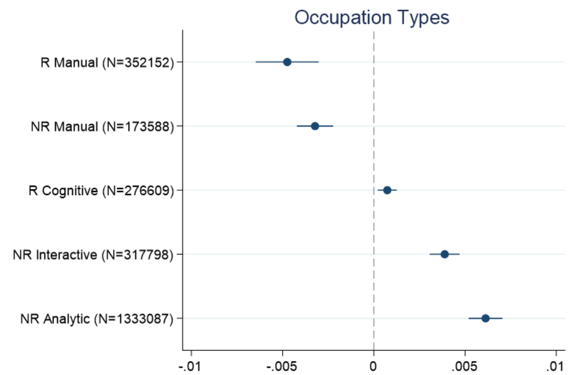

(b) Log Annual Earnings

NOTE. — All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, and 3-digit occupation FE. Earnings regressions also include year FE. Numerical results are reported in Appendix A6, Table E1. Point estimates are shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

pations, while workers in routine manual (–0.5%) and non-routine manual jobs (–0.3%) face losses.[18] Employment results, in turn, are small and inconsistent.[19]

Since AI mainly complements workers in cognitively demanding occupations, their payoff may depend on how specialized the occupational environment already is. If AI enhances the value of highly specialized expertise, workers in occupations with great emphasis on analytic or interactive tasks should benefit the most. If, instead, AI diminishes the value of expertise, some tasks may become substitutable and erode the advantage of expert workers. To assess this interplay, we estimate a triple-interaction model that combines workers' skill groups with dummies equal to one for occupations in the top tercile of the respective task distribution.

Figure 5 confirms that experts consistently realize the highest AI-induced earnings gains across all task groups, with total effects ranging from 1% to 2.2%. Yet, experts in occupations

[18]Contrary to our results on skill groups, earnings estimates in cognitive-intensive occupations are mainly driven by demand for developer-stage skills (Appendix A6, Table E5), suggesting worker skill groups and occupational task structure identify distinct domains of AI exposure.

[19]We find qualitatively similar results using task structure data from the German Occupation Panel (Grienberger, Janser, and Lehmer 2023), which is also based on core requirements in BERUFENET (Table E5). In Appendix A5 we further show that our vacancy-based task measures correlate highly with this data.

Figure 5: IV Triple Interaction Regressions of Annual Earnings on AI Exposure: By Worker Skill Groups and Occupational Task Intensity

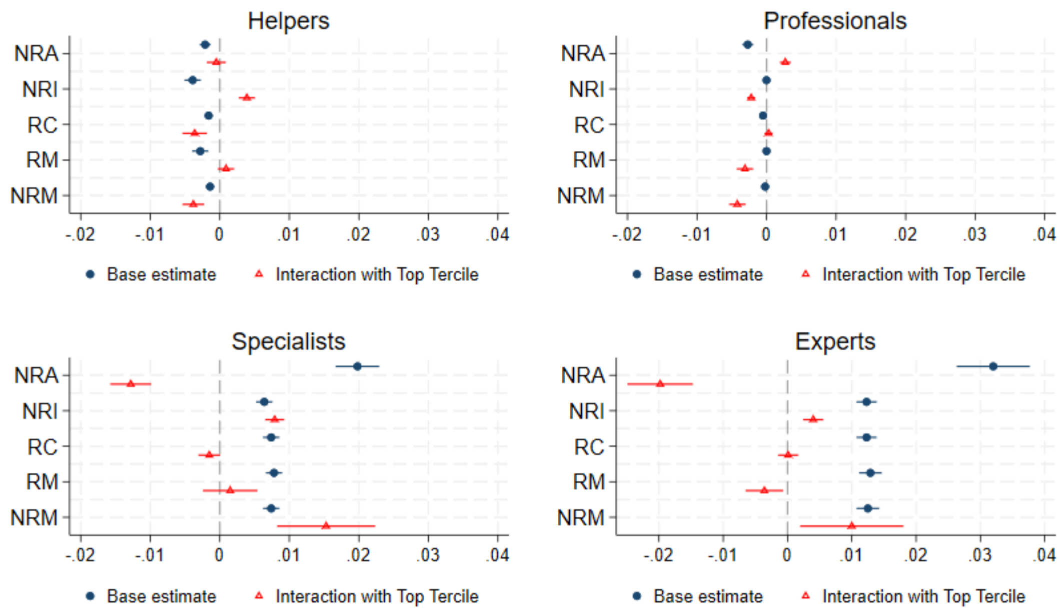

NOTE. — N = 2,385,502. All regressions are based on equation (4), with an additional interaction with a dummy = 1 if a 3-digit-occupation is in the top tercile of the underlying task distribution. We use controls described in Section 5, including worker, CZ, and 3d-occupation FE. Numerical results are reported in Appendix A6, Table E8. Point estimates are shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

where analytic tasks are most intensively applied do not capture the largest returns (e.g., software developers). A large base effect (+3%) is partly offset by a negative interaction effect (–2%), suggesting that AI partly substitutes for analytic tasks in the most analytic-intensive occupations. Instead, returns are stronger for experts in cognitive-intensive occupations with a more balanced task composition (e.g., occupations in accounting, controlling, and auditing) and interactive occupations (e.g., sales).[20]

Taken together, these results reaffirm that expertise is the decisive margin for identifying who benefits from AI, though the occupational task mix shapes the size of these payoffs. The interaction of worker skill groups with occupational task structures thus provides new insights on "winners" and "losers" from AI. However, this analysis remains silent on the role

[20]We also examined if occupational task structures shape employment stability across skill groups but found no meaningful patterns. Results are available upon request.

of new tasks that cannot yet be automated —an open question we turn to next.

7.2 Reinstatement Effect

Our third prediction states that new technologies create new tasks, thus extending the range of production-related activities (Acemoglu and Restrepo 2019; Autor, Chin, Salomons, and Seegmiller 2024). This reinstatement effect suggests complementarities between AI and (skilled) labor and could in principle amplify positive outcomes for high-skilled workers.

Construction of Task Space. To assess reinstatement effects, we track how workers' task spaces evolve over time. Building upon BERUFENET-based measures, we find on average eight non-AI skills in online job ads, with 40% analytic, 23% routine cognitive, 17% interactive, and the two manual groups together about 20%. Following AAHR, we classify skills as "emerging" or "fading" based on changes in relative frequency between 2017 and 2023, formally defined as positive skill change (PSC) and negative skill change (NSC):[21]

$$\mathrm{PSC}^{(s)}_{lo,2017\rightarrow 2023} = \sum_{k\in s} \max\left\{\left(\frac{\mathrm{kw}_{lo,2023}}{v^{\mathrm{ALL}}_{lo,2023}} - \frac{\mathrm{kw}_{lo,2017}}{v^{\mathrm{ALL}}_{lo,2017}}\right), 0\right\} \tag{5a}$$

$$\mathrm{NSC}^{(s)}_{lo,2017\rightarrow 2023} = -\sum_{k\in s} \min\left\{\left(\frac{\mathrm{kw}_{lo,2023}}{v^{\mathrm{ALL}}_{lo,2023}} - \frac{\mathrm{kw}_{lo,2017}}{v^{\mathrm{ALL}}_{lo,2017}}\right), 0\right\} \tag{5b}$$

where $\mathrm{kw}_{lo,t}$ is the number of times a keyword k belonging to task group s appears in LLM lo in years $t_0 = 2017$ and $t_1 = 2023$. Normalizing by the number of vacancies v^{ALL} in the same LLM captures the relative importance of each keyword at the start and end of the period. Around 45% of all skills display positive skill change (3.5 out of 7.7), the remainder displaying negative skill change. The resulting net skill change measure SC is then defined as the difference between annual demand for emerging and fading skills:

$$SC^{(s)}_{lot} = PSC^{(s)}_{lot} - NSC^{(s)}_{lot} \tag{6}$$

[21] See Deming and Noray (2020) and Peede and Stops (2024) for similar approaches.

Since both $PSC^{(s)}$ and $NSC^{(s)}$ are non-negative, $SC^{(s)}$ captures whether emerging skills outweigh fading skills in a given year. Variation in $SC^{(s)}$ within a worker's LLM over time thus proxies changes in the worker's relevant task space.

Results. Next, we explore if AI exposure is associated with task shifts. To this end, we estimate a modified version of our main IV regression eq. (3) at the LLM-year level, regressing the net skill change $SC^{(s)}$ of each skill group s on the LLM-specific AI exposure:

$$SC^{(s)}_{lot} = \phi_{lo} + \beta_{1,s}\widehat{AI}_{lot} + \beta_2 X_{lot} + \theta_t + \epsilon_{lot} \tag{7}$$

Figure 6 shows that an increase in AI exposure is associated with an expansion of the task space, though with discernible heterogeneity. A doubling in the share of AI vacancies expands the analytic task space by 4.7 pp., which corresponds to nearly one-and-a-half times the baseline average of 3.2 analytic requirements per job ad. The interactive task space also expands, though more modestly (+0.7 pp., baseline: 1.3). In contrast, both the non-routine manual (-1.5 pp., baseline: 0.6) and routine manual (-0.9 pp., baseline: 0.9) task spaces shrink. Echoing Gathmann, Grimm, and Winkler (2024), our results thus point to an AI-induced shift towards cognitively demanding tasks, though more strongly towards non-routine activities rather than "higher-order" routine cognitive ones as in their study.

To examine how these shifts translate into labor market outcomes, we estimate a modified version of our baseline IV regression (3):

$$\begin{aligned} Y_{ilot} = \alpha_i + \beta_1 \cdot \widehat{AI}_{lot} + \sum_{s=1}^{S} \gamma_s \cdot SC^{(s)}_{lot} + \sum_{s=1}^{S} \delta_s \cdot \widehat{AI}_{lot} \cdot SC^{(s)}_{lot} \\ + \beta_2 X_{it} + \psi_l + \omega_o + \theta_t + \epsilon_{ilot} \end{aligned} \tag{8}$$

We interact AI exposure with $SC^{(s)}$, such that the combined effect of AI on outcomes reflects both the magnitude of task shifts (i.e. the reinstatement effect) and their associated earnings or employment implications. Our parameters of interest are δ_s, capturing how much

Figure 6: Reinstatement Effect: IV Regressions of Net (Non-AI) Skill Change on AI Skill Demand, by Skill Group

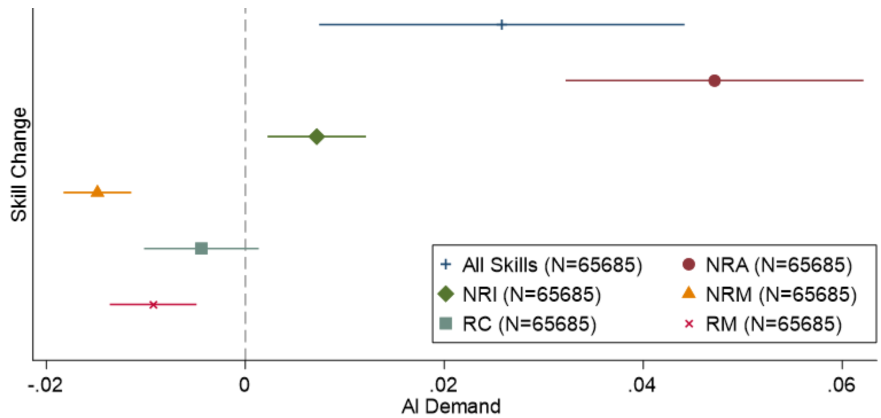

NOTE. — All regressions are based on equation (7) and include controls described in Section 5 in terms of LLM-shares. Numerical results are reported in Appendix A6, Table E9. Point estimates are shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

an AI-induced change in the task space translates into earnings and employment responses.[22]

Figure 7 shows that the results of our simple baseline analysis (column 4 in Table 3) are amplified once reinstatement effects are considered (+0.5 days employed, +0.8% earnings). The strongest contribution comes from analytic tasks. An AI-induced increase in the analytic task space by 1 pp. is associated with earnings gains of 0.17 %. Since a doubling in AI exposure expands this task space by roughly 4.7 pp., the total implied effect amounts to about 0.8%. Interactive tasks show a similar but smaller pattern, with a total estimate of 0.1% (0.7 pp. $\times$ 0.12 %). Routine cognitive activities display the strongest interaction coefficient (0.47 %), though their overall impact is muted because underlying task shifts are small and imprecisely estimated. The only exception is routine manual work. Here, AI exposure reduces the underlying task space, and since additional routine tasks tend to lower wages, this reduction actually results in a modest net earnings gain of +0.33%. Employment results reveal mostly insignificant estimates with small point estimates, suggesting no major employment implications of AI-induced reinstatement effects.

[22]Note we exclude 2-digit occupation-by-year FE to preserve meaningful variation in this demanding specification (though including these FE does not affect the results substantially).

Figure 7: IV Regressions of Annual Earnings and Employment Stability on AI Skill Demand and Net (Non-AI) Skill Change (Implications induces by Reinstatement Effect)

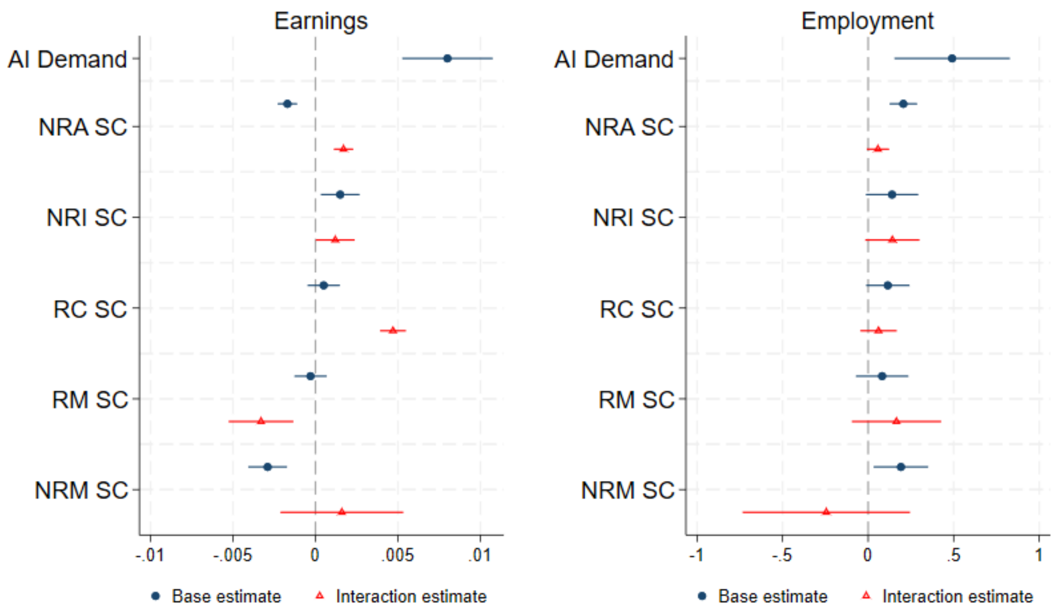

NOTE. — N = 2,078,928. All regressions are based on eq. (8) and controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. For earnings regressions we also include year FE. Numerical results are reported in Appendix A6, Table E10. Point estimates shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data); SIABv7523; own calculations.

Combined, AI-induced task shifts reinforce labor market gains primarily in analytic and interactive domains, while offering limited scope for the creation of new tasks in manual occupations. This pattern supports the interpretation of reinstatement effects as an additional channel through which workers benefit from AI via the expansion of cognitive tasks.

8 Discussion: Implications of AI on Inequality

Our analysis shows that high-skilled workers benefit in terms of employment and earnings, while lesser-skilled workers face higher displacement risks. These findings cast doubt on optimistic views of AI as a force for reducing inequality (Autor 2024). We thus conclude our analysis with a discussion of AI-induced implications on inequality. To this end, we assign workers to deciles of the earnings distribution (Table E11 in Appendix A6) and estimate

Figure 8: IV Regressions of Earnings and Employment on AI Skill Demand: By Income Bins

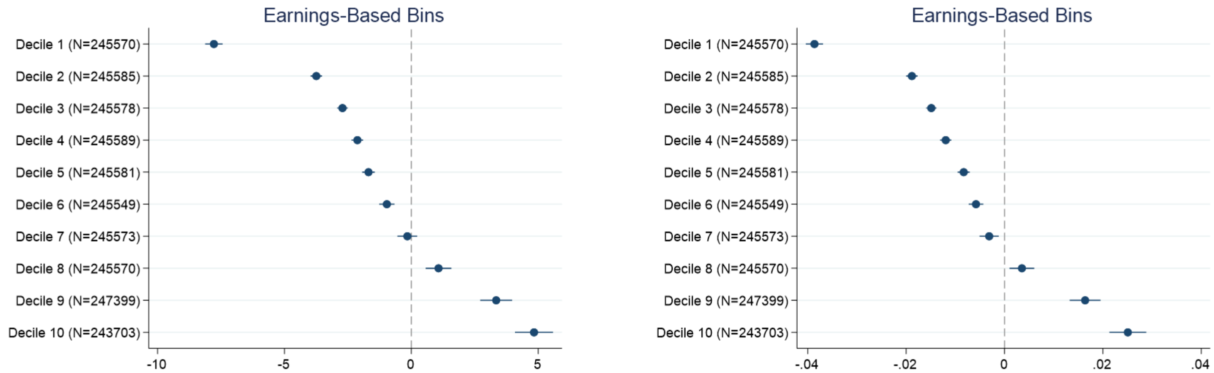

(a) Number of Days Employed

(b) Log Annual Earnings

NOTE. — All regressions are based on eq. (4) and include controls described in Section 5, including worker, CZ, 3d-occupation, and 2d-occupation-by-year FE. Earnings regressions also include year FE. Numerical results are reported in Appendix A6, Table E12. Point estimates are shown with 95% confidence intervals. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

group-specific implications by interacting AI exposure with bin dummies (see eq. 4).

Figure 8 reveals a steep gradient: workers in the bottom deciles experience significant losses, while those at the top gain. In the lowest decile, a doubling in the share of AI vacancies reduces employment by eight days and earnings by 3.9%. Negative results persist but weaken toward the middle of the distribution. From the eighth decile upward, estimates turn positive, reaching five additional days worked and a 2.5% earnings increase at the top. We caution that this analysis offers only suggestive evidence. Still, our findings echo concerns that AI may amplify existing inequalities in labor markets (Acemoglu 2025).

Viewed through the AAHR model, these results imply that productivity-enhancing effects of AI accrue mainly to high-wage workers. Our evidence aligns with findings of positive implications of AI for top earners and STEM workers, but displacement among lesser-skilled and production workers (Bonfiglioli, Crinò, Gancia, and Papadakis 2025). Conceptually, this divergence may reflect exposure to different technologies, for example if lesser-skilled workers are more exposed to automation technologies, but high-skilled workers to augmentation technologies (Autor, Chin, Salomons, and Seegmiller 2024; Marguerit 2025), though this distinction is beyond the scope of this paper. In terms of recent interest on productivity effects of (Gen)AI, our results suggest that high-productivity workers have gained most from (pre-Gen)AI technologies.[23]

9 Conclusions

We study how AI skill demand affects workers' employment stability and earnings outcomes and identify the channels shaping these responses. Using German online job vacancy data merged with administrative records for 2017–2023, we capture AI exposure with the share of vacancies requiring AI skills at the occupation-region-year level. Guided by the Acemoglu, Autor, Hazell, and Restrepo (2022) framework, we assess the relative importance of displacement and productivity effects by estimating the impact of changes in AI exposure on workers' outcomes. To address endogeneity from non-random exposure, we instrument local AI demand with national demand trends that are plausibly orthogonal to local market conditions (Azar, Marinescu, and Steinbaum 2022).

Our main result is that AI has not yet resulted in meaningful displacement or productivity effects on average. While we find associations with employment stability and earnings in simple baseline specifications, these vanish once we account for occupation-specific demand shocks in our preferred IV specification. However, these average results mask heterogeneity across skill groups. Using codified measures of job complexity —which distinguish between helpers, professionals, specialists, and experts —we identify domain-specific expertise as the strongest predictor of AI-induced gains. Expert workers consistently realize the largest AI-induced earnings gains, while non-experts often face losses, with the size of these gains shaped by occupational task structures.

[23] Stronger productivity boosts from GenAI tools for high-skilled workers have been shown in Otis, Clarke, Delecourt, Holtz, and Koning (2024) and Roldán-Monés (2025), and for lesser-skilled workers, by domain, for writing (Noy and Zhang 2023), advertising (Chen and Chan 2024), programming (Peng, Kalliamvakou, Cihon, and Demirer 2023), legal analysis (Choi, Monahan, and Schwarcz 2024), consulting (Dell'Acqua, McFowland III, Mollick, Lifshitz-Assaf, Kellogg, Rajendran, Krayer, Candelon, and Lakhani 2023), and customer support (Brynjolfsson, Li, and Raymond 2025).

Beyond skill heterogeneity, we provide worker-level evidence on the reinstatement effects of AI. We show that AI expands the task space of analytic and interactive activities while shrinking manual ones, and that these expansions are associated with measurable earnings gains, though no meaningful employment stability. This finding suggests that AI not only substitutes for existing tasks or raises productivity within them, but also creates new tasks, offering an additional channel through which workers can benefit from the emergence of AI.

Our results provide two key policy implications. First, identifying early "winners" and "losers" of AI is crucial for designing effective support mechanisms for vulnerable groups such as lesser-skilled workers, as our findings raise concerns about widening inequality in the early stages of AI (Acemoglu 2025). Second, the central role of expertise underscores the growing importance of domain knowledge and adaptable skills (Autor 2024; Autor and Thompson 2025). As AI adoption progresses, firms may increasingly shift toward skills-based hiring, raising important questions about how education and training systems should evolve to equip workers for AI-integrated jobs.

References

Acemoglu, Daron. 2002. Directed Technical Change. *The Review of Economic Studies* 69, no. 4: 781–809.

Acemoglu, Daron. 2025. The Simple Macroeconomics of AI. *Economic Policy* 40, no. 121: 13–58.

Acemoglu, Daron, Gary W. Anderson, David N. Beede, Catherine Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, et al. 2025. Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey. In *Technology, Productivity, and Economic Growth*, 13–55. Chicago: University of Chicago Press.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics* 40, no. S1: 293–340.

Acemoglu, Daron and Pascual Restrepo. 2019. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives* 33, no. 2: 3–30.

Acemoglu, Daron and Pascual Restrepo. 2020. Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy* 128, no. 6: 2188–2244.

Acemoglu, Daron and Pascual Restrepo. 2022. Tasks, Automation, and the Rise in US Wage Inequality. *Econometrica* 90, no. 5: 1973–2016.

Aghion, Philippe, Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard. 2025. How Different Uses of AI Shape Labor Demand: Evidence from France. *AEA Papers and Proceedings* 115: 62–67.

Albanesi, Stefania, António Dias da Silva, Juan F Jimeno, Ana Lamo, and Alena Wabitsch. 2024. New Technologies and Jobs in Europe. *Economic Policy* 40, no. 121: 71–139.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska. 2021. The demand for AI skills in the labor market. *Labour Economics* 71: 102002.

Arntz, M., M. Baum, E. Brüll, R. Dorau, M. Hartwig, F. Lehmer, B. Matthes, S.-C. Meyer, O. Schlenker, A. Tisch, and S. Wischniewski (2024a). *Digitalisierung und Wandel der Beschäftigung (DiWaBe 2.0): Eine Datengrundlage für die Erforschung von Künstlicher Intelligenz und anderer Technologien in der Arbeitswelt.* 1. Auflage. baua: Bericht. Dortmund: Bundesanstalt für Arbeitsschutz und Arbeitsmedizin.

Arntz, Melanie, Sabrina Genz, Terry Gregory, Florian Lehmer, and Ulrich Zierahn-Weilage (2024b). *De-Routinization in the Fourth Industrial Revolution - Firm-Level Evidence.* IZA Discussion Paper 16740. IZA - Institute of Labor Economics.

Aum, Sangmin and Yongseok Shin (2025). *The Labor Market Impact of Digital Technologies.* NBER Working Paper 33469. Cambridge, MA: National Bureau of Economic Research.

Autor, David (2024). *Applying AI to Rebuild Middle Class Jobs.* NBER Working Paper 32140. Cambridge, MA: National Bureau of Economic Research.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. New Frontiers: The Origins and Content of New Work, 1940–2018. *The Quarterly Journal of Economics* 139, no. 3: 1399–1465.

Autor, David and Neil Thompson. 2025. Expertise. *Journal of the European Economic Association* 23, no. 4: 1203–1271.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review* 103, no. 6: 2121–2168.

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song. 2014. Trade Adjustment: Worker-Level Evidence. *The Quarterly Journal of Economics* 129, no. 4: 1799–1860.

Azar, José, Ioana Marinescu, and Marshall Steinbaum. 2022. Labor Market Concentration. *Journal of Human Resources* 57, no. S: 167–199.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics* 151: 103745.

Babina, Tania, Anastassia Fedyk, Alex X. He, and James Hodson (June 2023). *Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition.* Working Paper 31325. National Bureau of Economic Research.

Bassier, Ihsaan, Alan Manning, and Barbara Petrongolo. 2025. Vacancy Duration and Wages. *The Review of Economics and Statistics*: 1–28.

Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia, and Ioannis Papadakis. 2025. Artificial Intelligence and Jobs: Evidence from US Commuting Zones. *Economic Policy* 40, no. 121: 145–194.

Bonney, Kathryn, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, and Keith Savage (2024). *Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey.* NBER Working Paper 32319. Cambridge, MA: National Bureau of Economic Research.

Börschlein, Erik-Benjamin, Mario Bossler, and Martin Popp (2024). *Scarce Workers, High Wages?* IZA Discussion Paper 17447. IZA - Institute of Labor Economics.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. Quasi-Experimental Shift-Share Research Designs. *Review of Economic Studies* 89, no. 1: 181–213.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. A Practical Guide to Shift-Share Instruments. *Journal of Economic Perspectives* 39, no. 1: 181–204.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. Generative AI at Work. *The Quarterly Journal of Economics* 140, no. 2: 889–942.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings* 108: 43–47.

Büchel, Jan, Vera Demary, Henry Goecke, Enno Kohlisch, Oliver Koppel, Armin Mertens, Christian Rusche, Marc Scheufen, and Jan Wendt (2021). *KI-Monitor 2021: Status quo der Künstlichen Intelligenz in Deutschland.* Gutachten im Auftrag des Bundesverbandes Digitale Wirtschaft (BVDW) e.V. Institut der deutschen Wirtschaft.

Chen, Zenan and Jason Chan. 2024. Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise. *Management Science* 70, no. 12: 9101–9117.

Choi, Jonathan H, Amy B Monahan, and Daniel Schwarcz. 2024. Lawyering in the age of artificial intelligence. *Minnesota Law Review* 109: 147–218.

Dauth, Wolfgang and Johann Eppelsheimer. 2020. Preparing the sample of integrated labour market biographies (SIAB) for scientific analysis: a guide. *Journal for Labour Market Research* 54, no. 1: 10.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum. 2021. Adjusting to Globalization in Germany. *Journal of Labor Economics* 39, no. 1: 263–302.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner. 2021. The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association* 19, no. 6: 3104–3153.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani (2023). *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*. Harvard Business School: Technology & Operations Management Unit Working Paper 24-013. Harvard Business School.

Deming, David J. and Kadeem Noray. 2020. Earnings Dynamics, Changing Job Skills, and STEM Careers. *The Quarterly Journal of Economics* 135, no. 4: 1965–2005.

Dengler, Katharina, Britta Matthes, and Wiebke Paulus (2014). *Occupational Tasks in the German Labour Market: An alternative measurement on the basis of an expert database*. FDZ Methodology Report 12/2014. Research Data Centre (FDZ) of the Institute for Employment Research (IAB).

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In: *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Vol. 1, pp. 4171–4186.

Engberg, E., M. Hellsten, F. Javed, M. Lodefalk, R. Sabolová, S. Schroeder, and A. Tang. 2025. Artificial Intelligence, Hiring and Employment: Job Postings Evidence from Sweden. *Applied Economics Letters*: 1–6.

Engberg, Erik, Holger Görg, Magnus Lodefalk, Farrukh Javed, Martin Längkvist, Natália Pimenta Monteiro, Hildegunn Kyvik Nordås, Sarah Schroeder, and Aili Tang (2024). *AI Unboxed and Jobs: A Novel Measure and Firm-Level Evidence from Three Countries*. IZA Discussion Paper 16717. IZA - Institute of Labor Economics.

Eurostat (2024). *Use of Artificial Intelligence in Enterprises*. https://ec.europa.eu/eurostat/statistics-explained/index.php/Use_of_artificial_intelligence_in_enterprises. Accessed: 2025-09-15. European Commision.

Falck, Oliver, Anna Kerkhof, and Anita Wölfl. 2024. Künstliche Intelligenz – wie Unternehmen sie nutzen und was sie noch daran hindert. *ifo Schnelldienst* 77, no. 09: 57–63.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42, no. 12: 2195–2217.

Felten, Edward W, Manav Raj, and Robert Seamans (2023). *Occupational Heterogeneity in Exposure to Generative AI*. SSRN Pre-Print 4414065. Social Science Research Network.

Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter. 2006. Imputation Rules to Improve the Education Variable in the IAB Employment Subsample. *Journal of Contextual Economics* 126, no. 3: 405–436.

Fossen, Frank M. and Alina Sorgner. 2022. New digital technologies and heterogeneous wage and employment dynamics in the United States: Evidence from individual-level data. *Technological Forecasting and Social Change* 175: 121381.

Gathmann, Christina and Felix Grimm (2022). *The Diffusion of Digital Technologies and its Consequences in the Labor Market. Beiträge zur Jahrestagung des Vereins für Socialpolitik*

2022: Big Data in Economics. Conference Paper. ZBW - Leibniz Information Centre for Economics.
Gathmann, Christina, Felix Grimm, and Erwin Winkler (2024). *AI, Task Changes in Jobs, and Worker Reallocation*. IZA Discussion Paper 17554. IZA - Institute of Labor Economics.
Genz, Sabrina, Terry Gregory, Markus Janser, Florian Lehmer, and Britta Matthes (2021). *How Do Workers Adjust When Firms Adopt New Technologies?* IZA Discussion Paper 14626. IZA - Institute of Labor Economics.
Giering, Oliver, Alexandra Fedorets, Jule Adriaans, and Stefan Kirchner. 2021. Künstliche Intelligenz in Deutschland: Erwerbstätige wissen oft nicht, dass sie mit KI-basierten Systemen arbeiten. *DIW Wochenbericht* 48: 783–789.
Graf, Tobias, Stephan Grießemer, Markus Köhler, Claudia Lehnert, Andreas Moczall, Martina Oertel, Alexandra Schmucker, Andreas Schneider, and Philipp vom Berge (2025). *Schwach anonymisierte Version der Stichprobe der Integrierten Arbeitsmarktbiografien (SIAB) – Version 7523 v1*. Research Data Center (FDZ) of the Institute for Employment Research (IAB).
Grienberger, Katharina, Markus Janser, and Florian Lehmer. 2023. The Occupational Panel for Germany. *Journal of Economics and Statistics* 243, no. 6: 711–724.
Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller (2025). *Artificial Intelligence and the Labor Market*. Working Paper 33509. National Bureau of Economic Research.
Hunt, Jennifer, Iain M Cockburn, and James Bessen (2024). *Is Distance from Innovation a Barrier to the Adoption of Artificial Intelligence?* NBER Working Paper 33022. Cambridge, MA: National Bureau of Economic Research.
Kosfeld, Reinhold and Alexander Werner. 2012. Deutsche Arbeitsmarktregionen: Neuabgrenzung nach den Kreisgebietsreformen 2007–2011. *Spatial Research and Planning* 70, no. 1: 49–64.
Lewandowski, Piotr, Karol Madoń, and Albert Park (2025). *Workers' Exposure to AI Across Development Stages*. IZA Discussion Paper 18235. IZA - Institute of Labor Economics.
Lichter, Andreas, Andreas Peichl, and Sebastian Siegloch. 2015. The Own-Wage Elasticity of Labor Demand: A Meta-Regression Analysis. *European Economic Review* 80: 94–119.
Marguerit, David (2025). *Augmenting or Automating Labor? The Effect of AI Development on New Work, Employment, and Wages*. SSRN Pre-Print 5169611. Social Science Research Network.
McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. 2024. AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy* 33, no. 2: 375–415.
Noy, Shakked and Whitney Zhang. 2023. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science* 381, no. 6654: 187–192.
Otis, Nicholas, Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand Koning (2024). *The Uneven Impact of Generative AI on Entrepreneurial Performance*. SSRN Pre-Print 4671369. Social Science Research Network.
Ozgul, Pelin, Marie-Christine Fregin, Michael Stops, Simon Janssen, and Mark Levels (2024). *High-skilled Human Workers in Non-Routine Jobs are Susceptible to AI Automation but Wage Benefits Differ between Occupations*. arXiv Pre-Print 2404.06472. arXiv.

Peede, L. and M. Stops (2024). *Artificial Intelligence Technologies, Skills Demand and Employment: Evidence from Linked Job Ads Data*. IAB-Discussion Paper 15/2024. IAB - Institute for Employment Research.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer (2023). *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot*. arXiv Pre-Print 2302.06590. arXiv.

Prytkova, Ekaterina, Fabien Petit, Deyu Li, Sugat Chaturvedi, and Tommaso Ciarli (2024). *The Employment Impact of Emerging Digital Technologies*. CESifo Working Paper 10955. CESifo.

Rammer, Christian (2022). *Kompetenzen und Kooperationen zu Künstlicher Intelligenz: Ergebnisse einer Befragung von KI-aktiven Unternehmen in Deutschland*. Ergebnisse einer Befragung von KI-aktiven Unternehmen in Deutschland, Bundesministerium für Wirtschaft und Klimaschutz (BMWK). ZEW - Leibniz-Zentrum für Europäische Wirtschaftsforschung.

Rammer, Christian, Gastón P. Fernández, and Dirk Czarnitzki. 2022. Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Research Policy* 51, no. 7: 104555.

Roldán-Monés, Antonio (2025). *When GenAI increases inequality: Evidence from a university debating competition*. poid Working Paper 096. Program on Innovation and Diffusion.

Schmucker, Alexandra and Philipp vom Berge (2025). *Sample of Integrated Labour Market Biographies (SIAB) 1975 - 2023. FDZ Datenreport*. FDZ Data Report 02/2025. Research Data Centre (FDZ) of the Institute for Employment Research (IAB).

Stapleton, Katherine, Alexander Copestake, and Ashley Pople (2025). *AI and Services-Led Growth: Evidence from Indian Job Adverts*. SSRN Pre-Print 3957858. Social Science Research Network.

Taska, Bledi, Layla O'Kane, and Julia Nania (2022). *Artificial Intelligence in the UK: The relevance of AI in the digital transformation of the UK labour market*. Report. Lightcast.

Tolan, Songül, Annarosa Pesole, Fernando Martínez-Plumed, Enrique Fernández-Macías, José Hernández-Orallo, and Emilia Gómez. 2021. Measuring the Occupational Impact of AI: Tasks, Cognitive Abilities and AI Benchmarks. *Journal of Artificial Intelligence Research* 71: 191–236.

Webb, Michael (2020). *The Impact of Artificial Intelligence on the Labor Market*. SSRN Pre-Print 3482150. Social Science Research Network.

(Online) Appendix

AI in Demand:

How Expertise Shapes its (Early) Impact on Workers

A1 Details on OJV Data

A1.1 NLP Steps

Upon receiving the raw data from our data provider, Palturai/ Finbot, we link firm and vacancy information and preprocess the textual data, following conventions in the literature (Ash and Hansen 2023; Gentzkow, Kelly, and Taddy 2019). Beyond these basic steps, we also enrich the data. First, we assign each vacancy to a specific location, either at the zip code (39% of OJVs), municipality-level (48%), or county-level (10%).[1] Overall, 97% of postings can be linked to a county. Second, we classify job titles according to the German Classification of Occupations 2010 (KldB 2010v2020). Using codified job titles at the 8-digit level from the Federal Employment Agency, we directly assign about 60% of vacancies to "3-plus-5-digit" occupations, i.e. 3-digit occupations plus the fifth digit that captures differential job complexity within a 3-digit occupation (helpers, professionals, specialists, experts). For the remainder, we annotate unclassified postings.

A key step in preparing OJVs is the identification of distinct text segments ("Zoning").[2] Ideally, a posting contains four sections: (i) job activities, (ii) qualifications, (iii) benefits, and (iv) company description, though not all are always present. We implement zoning in two steps. First, we use ChatGPT-4o-mini to extract these sections from 50,000 postings, filtering out irrelevant text (e.g., HTML remnants, contact details). Second, we train a

[1] About 10% of postings lack detailed workplace information (often smaller firms with one location). In such cases, we use the imprint address. If multiple locations span several counties, postings are proportionally allocated.

[2] See the OJA Guide, developed by the Bertelsmann Stiftung and the Federal Institute for Vocational Education and Training, with Johannes Müller (&effect data solutions) as lead author.

German-language BERT model as a token classifier on this labeled set to predict which section each word belongs to. After manual validation, we apply the model to all postings.

Duplicate detection is also essential. While Finbot removes duplicates within platforms, postings can still appear across job posting sites. We identify duplicates when jobs are posted by the same firm, in the same location, for the same occupation, and within a 60-day window. We then compare textual similarity in the activity and qualification segments using pairwise partial Levenshtein distances. Based on a labeled set of 1,500 potential duplicates, this model performs best. Applying it, we exclude 13% of postings as duplicates.

Finally, we focus on information in the description of job activities and qualifications to identify job-specific skill requirements. A naive keyword-based approach is susceptible to false positives, so we adopt an n-gram approach, combining multiple search terms. In particular, job-specific skill requirements are identified whenever multiple search terms co-occur within a five-word window. For example, the term "support" itself is ambiguous, while the term "support clients" signals a relevant task in the context of interactive activities.

A1.2 External Validity

Online job vacancies (OJV) represent only one of many search channels, yet they are the dominant channel for recruiting high-skilled workers (Carrillo-Tudela, Kaas, and Lochner 2023) —including those with AI-related skills. While this concentration limits representativeness, it makes OJV data particularly suitable for identifying AI demand. Still, the non-representative nature of OJV data may distort results, raising concerns about generalizability.

We thus provide external validity using two sources: (i) the German Job Vacancy Survey (JVS), a representative survey conducted by the IAB (Bossler, Gürtzgen, Kubis, Küfner, and Popp 2021) and (ii) the number of workers subject to social security contributions (SSC) from official statistics of the German Federal Employment Agency (FEA). The JVS allows us to compare aggregate vacancy dynamics, while SSC data detailed benchmarks over time, occupations, and regions. Overall, our OJV data mirror JVS dynamics and yield similar stylized facts when reweighted by employment rather than the number of job postings.

I. Comparison with Job Vacancy Survey

Following common practice (Hershbein and Kahn 2018; Rengers 2018), we compare our OJV data with representative vacancy information. The JVS tracks labor demand and recruitment behavior since 1989 (Gürtzgen, Lochner, Pohlan, and van den Berg 2021).

Figure A1: Number of Vacancies over Time (2017–2024)

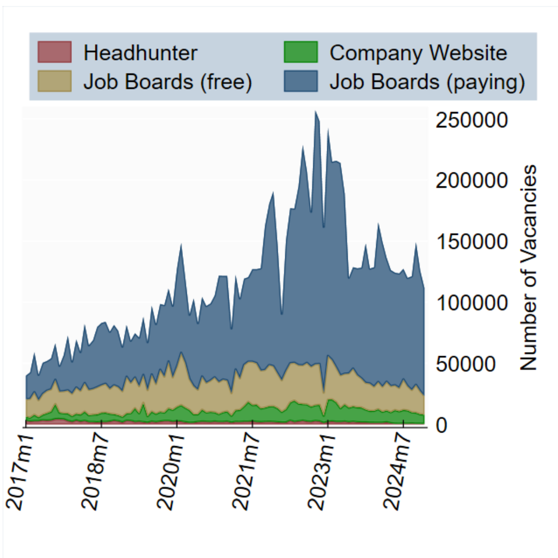

(a) OJV data, by source (Inflow)

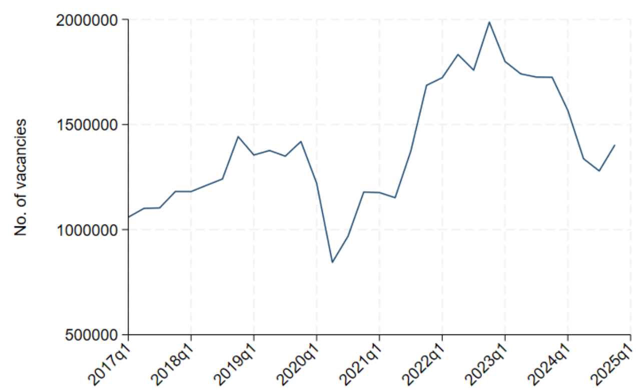

(b) IAB Job Vacancy Survey (Stock)

NOTE. —Panel A1a displays the number of online job vacancies that are posted each month in our data, i.e., monthly inflows, broken down by the type of source platform. Panel A1b displays the stock of vacancies firms taken from the IAB job vacancy survey by quarter. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; IAB Job Vacancy Survey; own calculations.

Figure A1 shows that OJV inflows between 2017–24 closely track vacancy stocks in the JVS. Panel A1a details OJV inflows by source platform, with job boards rising from 50% of postings in 2017 to 70% in 2023. Panel A1b provides fewer breakdowns but mirrors overall dynamics. Both datasets show an increasing trend of the number of postings over time, but with a sharp decrease at the beginning of the COVID-19 pandemic in 2020 and a subsequent rebound of postings. Specifically, JVS data record a 40% drop in vacancy stocks between 2019Q4 and 2020Q2, while OJV inflows fall by 30% from Dec 2019 to June 2020. These magnitudes match evidence from Australia (Shen and Taska 2020), Austria (Bamieh and Ziegler 2020; Bamieh and Ziegler 2022), Sweden (Hensvik, Le Barbanchon, and Rathelot 2021), the UK (Arthur 2021), and the US (Forsythe, Kahn, Lange, and Wiczer 2020). We thus conclude that our OJV data align with the cyclicality and shocks captured in representative surveys.

II. Comparison with number workers subject to SSC

Next, we test whether the skewness of our OJV data towards high-skilled jobs biases our stylized facts on AI demand. Our baseline AI measure weights by the number of job postings to account for differential job posting behavior among firms. To assess if this skewness distorts the validity of our measure we reweight it by SSC employment at the occupation–region–year level (Federal Employment Agency 2023).[3]

Figure A2: Number of online job vacancies over time (2017–2023)

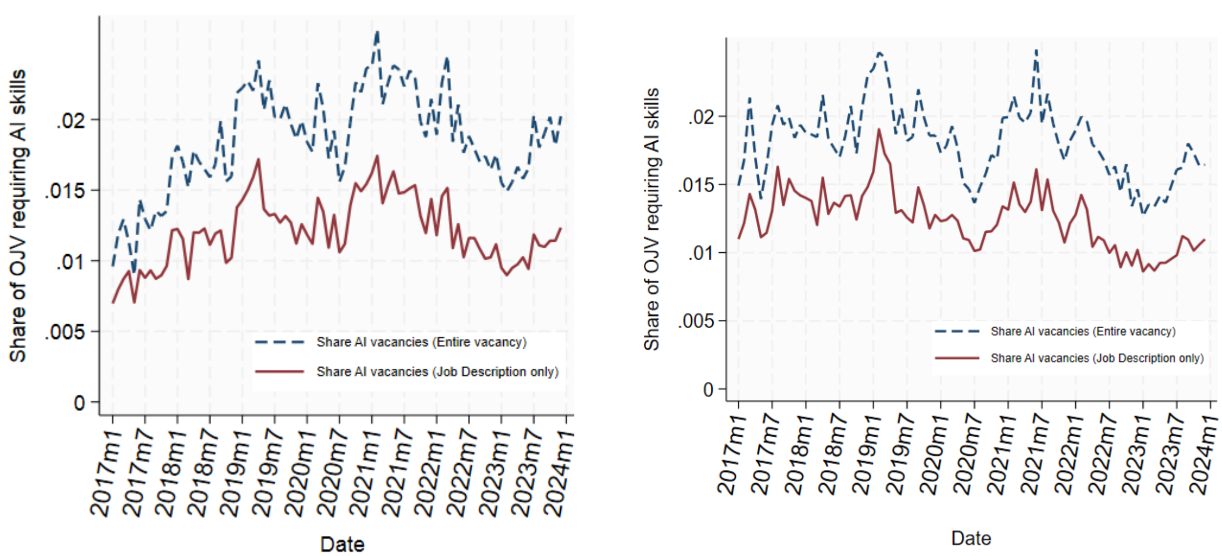

(a) Trends in AI Skill Demand: Weighting by number of postings

(b) Trends in AI Skill Demand: Weighting by number of SSC workers

NOTE. —Panel A2a shows our baseline AI skill demand, weighted by the number of postings in a local labor market. Panel A2b shows AI skill demand weighted by the number of workers subject to social security contributions. The red solid defines AI skills when extracted from the job profile (our baseline), while the blue dashed line uses the entire job ad text. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; Official Statistics of the German Federal Employment Agency 2017–2023; own calculations.

[3]Note that the JVS is less suited for this exercise due to limited occupational and regional detail.

Panel A2a shows the share of AI vacancies between 2017–2023 in our baseline specification, weighted by postings, while Panel A2b uses SSC worker weights. Both panels reveal very similar dynamics: a rise in AI demand until early 2020, a COVID-induced slowdown, a rebound in 2021, with a subsequent slowdown from 2022 onward. These common patterns suggest that the skew toward high-skilled postings does not distort our results meaningfully.

Figure A3 compares occupations by baseline AI demand in 2017 (horizontal axis) and subsequent changes through 2023 (vertical axis). While panel A3a shows that certain occupations with a large share of high-skilled workers are overrepresented in our OJV data compared to the SSC distribution (Panel A3b), both weighting schemes yield the same core insight: occupations with relatively high initial AI demand are also those that experience the strongest subsequent growth.

Finally, Figure A4 then illustrates regional patterns. Panels A4a and A4b show the share of AI vacancies across 141 CZs in 2017, weighted by postings and SSC workers respectively. Both panels indicate strong concentration in the South and West of Germany. Panels A4c and A4d display the same distributions in 2023 using the 2017 class boundaries, underlining the subsequent diffusion of AI demand across regions.

Figure A3: Dynamics in Occupational Demand for AI Skills (2017–2023)

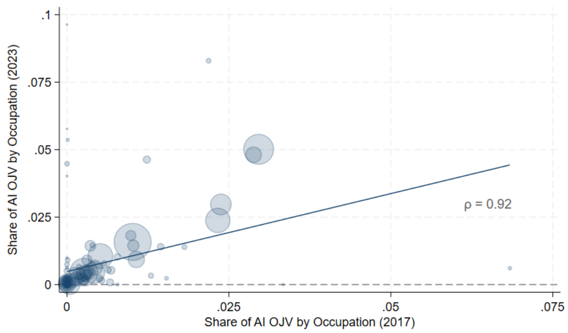

(a) OJV Weights

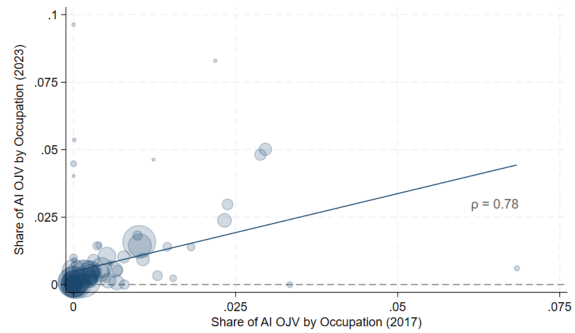

(b) SSC Weights

NOTE. —The X-axis displays the share of OJV with AI demand for 3-digit occupations as of 2017, while the Y-axis displays the change in AI Vacancies between 2017–2023. Panel A3a is weighted by the number of job postings, while Panel A3b is weighted by the number of workers who are subject to social security contributions (SSC) to enhance representativity of the data. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; Official Statistics of the German Federal Employment Agency; own calculations.

Figure A4: Demand for AI Skills in Germany Across Commuting Zones (2017–2023)

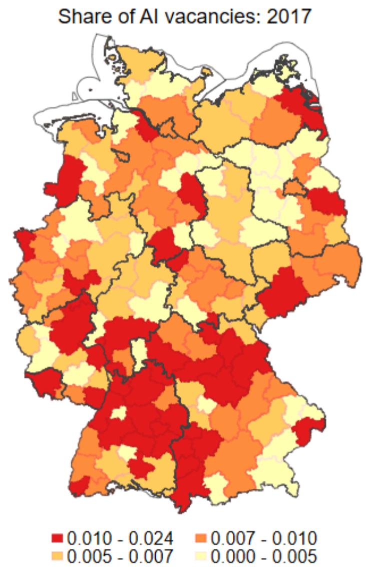

(a) AI Skill Demand: 2017 (OJV weights)

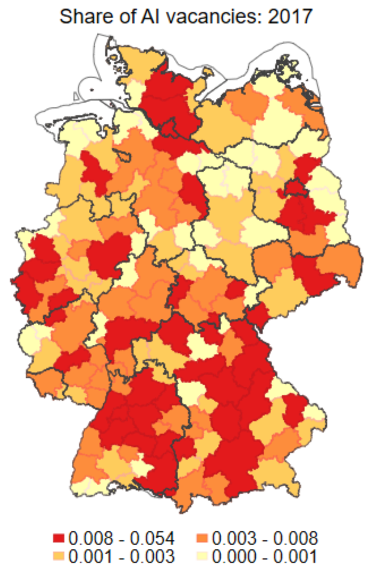

(b) AI Skill Demand: 2017 (SSC weights)

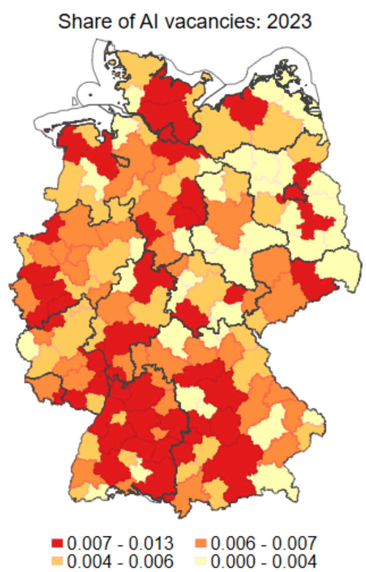

(c) AI Skill Demand: 2023 (OJV weights)

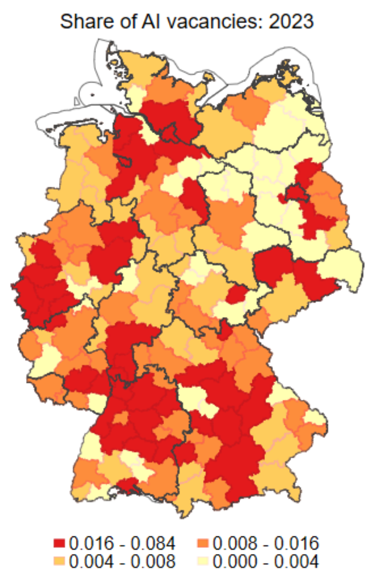

(d) AI Skill Demand: 2023 (SSC weights)

NOTE. —Local labor markets are assigned into four classes of task intensity. Each class corresponds to quartiles as of 2017 where lowest quartile implies lowest AI demand (yellow) and highest quartile implies highest AI demand (red). Panels A4a and A4c represent our baseline scenario, in which we weight AI skill demand by the number of job postings within a local labor market. In contrast, Panels A4b and A4d use the number of workers who are subject to social security contributions (SSC) as weights to enhance representativity of the data. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; Official Statistics of the German Federal Employment Agency; own calculations.

A2 Details on Construction of AI Measure

A2.1 Overview AI skills

In order to identify AI skills from online job vacancies, we employ a three-step procedure, involving (i) a keyword-based approach, (ii) manual annotation, and (iii) AI-assisted refinements with ChatGPT.

In the first step, we create a keyword list, especially the Lightcast taxonomy on AI (also used in Acemoglu, Autor, Hazell, and Restrepo (2022)), and supplemented it with German keywords from Büchel, Demary, Goecke, Kohlisch, Koppel, Mertens, Rusche, Scheufen, and Wendt (2021). This step is important, since some postings use German terms (e.g., maschinelles Lernen) instead of the more common English expressions (machine learning). While drawing on prior studies enriches our taxonomy, the preliminary list had clear limitations, especially (i) lack of coverage of recent innovations such as Large Language Models, (ii) industry-specific terminology relevant in the German context, and (iii) excluded abbreviations or colloquial phrasing common in postings.

In the second step, we address these shortcomings with an extended list through manual annotation. Two trained student assistants independently labeled AI skills in a sample of 600 postings, following strict annotation rules and predefined categories. Across five rounds, inter-annotator agreement ranged from 60% to 87%, with an average of 76%. To formally quantify consistency, we computed Cohen's Kappa, which yielded 0.41. This value falls in the range of moderate agreement (Landis and Koch 1977), a level common in classifications where definitions are fluid and still evolving. Disagreements, which often arose for ambiguous cases such as "automation software" or "big data analysis" were resolved by the authors.

Figure B1: Word Clouds of AI Skills: Baseline Definition

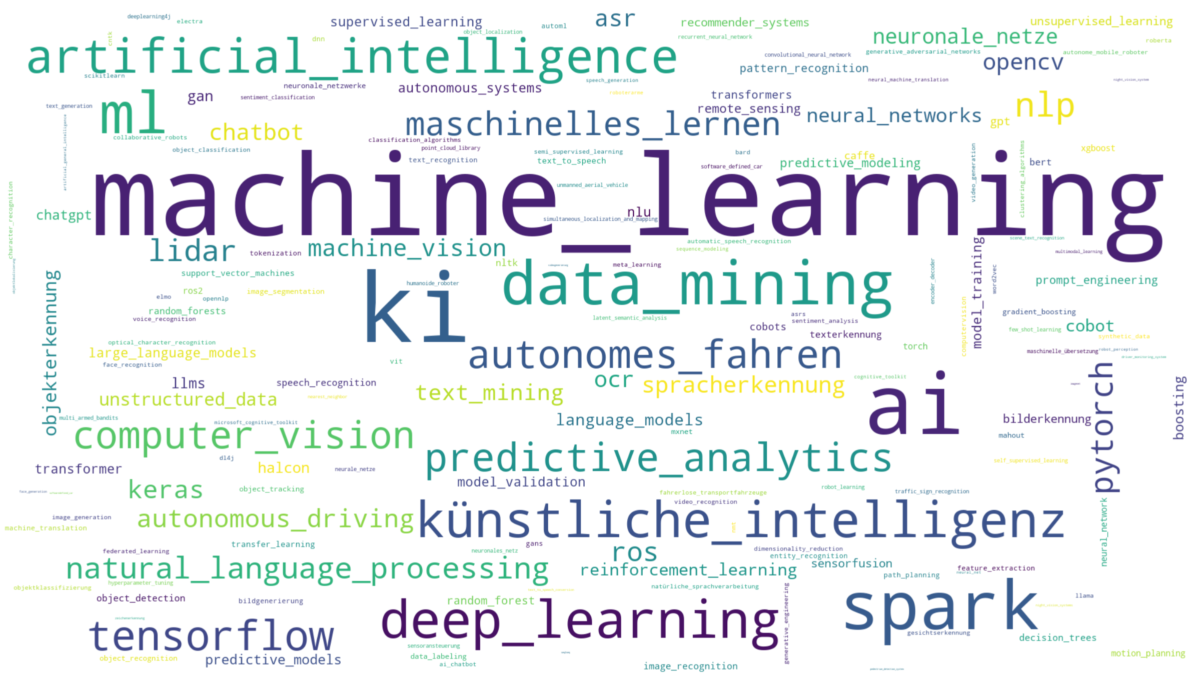

NOTE. —This word cloud comprises keywords that are associated with AI skills. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2024; own calculations.

In the third step, we incorporated AI-assisted refinements. Annotators used ChatGPT 4.0 to (i) classify ambiguous terms, (ii) validate emerging AI-related keywords, and (iii) support inclusion decisions for contested cases. ChatGPT was prompted with our taxonomy, classification rules, and distinctions between AI Methods, AI Applications, and Generative AI. For example, annotators asked: "In job postings, 'Digital Twin Technology' appears under required skills. Should this be classified as an AI skill? Provide reasoning and assess relevance on a scale from 1 (not AI-related) to 5 (highly AI-relevant)."

Table B1: Most Common AI-Related Keywords in Job Ads (Top 40)

Rank	Keyword	Share
1	machine learning	15.02%
2	ki	11.75%
3	ai	9.39%
4	künstliche intelligenz	5.42%
5	chatbot	4.76%
6	artificial intelligence	4.37%
7	autonomes fahren	3.96%
8	spark	3.77%
9	data mining	3.69%

Continued on next page

Table B1 – continued from previous page

Rank	Keyword	Share
10	ml	3.60%
11	adas	3.40%
12	deep learning	2.63%
13	text mining	2.37%
14	predictive analytics	2.08%
15	computer vision	1.74%
16	maschinelles lernen	1.27%
17	tensorflow	1.24%
18	lidar	1.18%
19	autonomous driving	1.04%
20	machine vision	0.95%
21	ros	0.95%
22	nlp	0.92%
23	natural language processing	0.78%
24	spracherkennung	0.68%
25	new mobility	0.58%
26	ocr	0.55%
27	boosting	0.54%
28	pytorch	0.52%
29	sw design	0.50%
30	remote sensing	0.47%
31	bert	0.47%
32	asr	0.39%
33	neural networks	0.38%
34	keras	0.36%
35	neuronale netze	0.35%
36	adtf	0.34%
37	objekterkennung	0.32%
38	opencv	0.31%
39	reinforcement learning	0.29%
40	v2x	0.29%

NOTE. — The share is computed relative to the total number of AI keyword mentions between 2017 and 2023. Spelling variants and synonyms (e.g., ”ki”, ”AI”) are listed separately. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; own calculations.

We adopted a conservative approach in the sense that only terms scored 4 or 5 were included in our keyword list. This procedure ensured that emerging concepts (e.g., “Neural Radiance Fields” or “Transformer Architectures”) were captured, while overly broad digital skills (e.g., business intelligence) were excluded. Final inclusion decisions were made by the authors, with ChatGPT serving as auxiliary guidance.

Overall, combining keyword-based identification, manual annotation, and AI-assisted refinements allows us to construct a comprehensive taxonomy consistent with prior work, but also adaptive to recent developments, especially in Generative AI. Figure B1 shows the resulting AI skills via a word cloud. Table B1 summarizes the 40 most frequent AI keywords, and Table B2 displays the Top 20 and Bottom 20 occupations by AI vacancy share.[4]

[4]As expected, occupations with high AI shares include IT-related and technical jobs, but also managerial

Table B2: AI Skill Demand by Occupations (Top 20 and Bottom 20)

Rank	Occupation *(3-digit KLdB 2010)*	AI Share	Job ads (in k)
	Panel A: Top 20 occupations		
1	Philology *(911)*	28.59%	0.45
2	Computer science *(431)*	10.52%	308.6
3	Teachers/researchers at universities *(843)*	8.12%	54.1
4	Mathematics and statistics *(411)*	7.55%	8.2
5	Aircraft pilots *(523)*	5.68%	0.6
6	IT-system-analysis, IT-app. & IT-sales *(432)*	5.40%	127.8
7	Software development and programming *(434)*	5.13%	442.0
8	Physics *(414)*	3.49%	9.8
9	Product and industrial design *(931)*	3.04%	21.9
10	IT-network eng., IT-coordination, IT-admin. *(433)*	3.03%	275.5
11	Technical R&D *(271)*	2.98%	247.1
12	Economics *(914)*	2.96%	2.1
13	Police and criminal investigation, jurisdiction *(532)*	2.12%	1.4
14	Insurance and financial services *(721)*	1.95%	81.1
15	Business organisation and strategy *(713)*	1.70%	886.8
16	Social sciences *(913)*	1.65%	14.9
17	Technical media design *(232)*	1.59%	52.0
18	Public relations *(922)*	1.48%	8.3
19	Surveying and cartography *(312)*	1.29%	8.3
20	Laboratory occupations in medicine *(812)*	1.26%	29.9
	Panel B: Bottom 20 occupations		
112	Paper-making and -processing, packaging *(231)*	0.00%	3.3
113	Construction and transport equipment ops. *(525)*	0.00%	23.8
114	Occup. health, safety admin, disinfection *(533)*	0.00%	3.8
115	Musical instrument making *(936)*	0.00%	0.1
116	Stage, costume and prop design *(946)*	0.00%	0.6
117	Sales (retail) books, antiques, music *(625)*	0.00%	0.7
118	Funeral services *(824)*	0.00%	0.3
119	Floor laying *(331)*	0.00%	5.2
120	Sales (retail) durables (e.g. clothing, cars) *(622)*	0.00%	20.9
121	Environmental protection engineering *(422)*	0.00%	0.9
122	Sales (retail) foodstuffs *(623)*	0.00%	18.0
123	Ceramics and glassware (artisanal) *(934)*	0.00%	0.1
124	Body care services *(823)*	0.00%	8.8
125	Performing arts *(942)*	0.00%	0.2
126	Industrial glass-making *(213)*	0.00%	1.2
127	Ship's officers and masters *(524)*	0.00%	0.4
128	Artisan craftwork and fine arts *(933)*	0.00%	0.3
129	Passenger transport services *(514)*	0.00%	3.3
130	Industrial ceramic-making *(214)*	0.00%	0.2
131	Theology and church services *(833)*	0.00%	0.3

NOTE. — This table presents the Top 20 and Bottom 20 occupations (3-digit KLdB 2010 level) in terms of their share of Online Job Vacancies (OJV) with AI-related skill demand. Job ads refer to the total number of job ads by occupation, between 2017 and 2023. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; own calculations.

positions. Some less obvious occupations appear as well, such as Philology (linked to machine translation and natural language applications). Since this occupation comprises only 450 postings in our sample, however, even a small number of AI-tagged vacancies produces a high share.

A2.2 AI measure: Internal Validity

To provide deeper insights into distinct AI domains, we divide our baseline measure into two subcategories. The first category —AI methods —reflects the developer-perspective, comprising algorithms, methods, and software tools such as TensorFlow, PyTorch, support vector machines, and Transformer models. The second category —AI applications —reflects the user-perspective, capturing industry-specific use cases such as autonomous driving, LIDAR, and AI-powered fraud detection in finance.

We use ChatGPT 4.0 to classify skills into these two categories, providing contextual information on our taxonomy and asking whether terms should be categorized as methods or applications. This process ensures consistency and helps capture emerging AI terminology. We have used the following prompt to classify AI skills with ChatGPT 4.0:

"*I have a list of AI-related skills and need to classify each into one of the two categories described below.*

AI Methods represent the AI development perspective and include general algorithms, methods, and software tools such as TensorFlow, PyTorch, support vector machines, and non-generative Transformer models.

AI Applications represent the AI user perspective and include industry-specific use cases where AI tools are applied to real-world settings. Examples include AI-powered fraud detection in finance, AI in autonomous driving such as LIDAR and ADAS, and AI applications in retail, manufacturing, or supply chain automation.

Classification follows the following criteria: If a skill relates to AI Methods, it should be classified as AI Methods. If a skill refers to practical AI-powered use cases in specific industries, it should be classified as AI Applications."

Figure B2: Word Clouds of AI Skills by Category

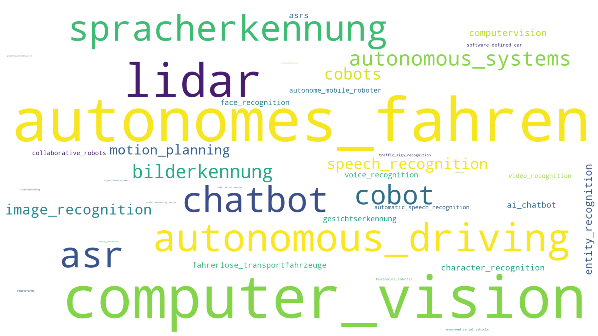

(a) AI Applications

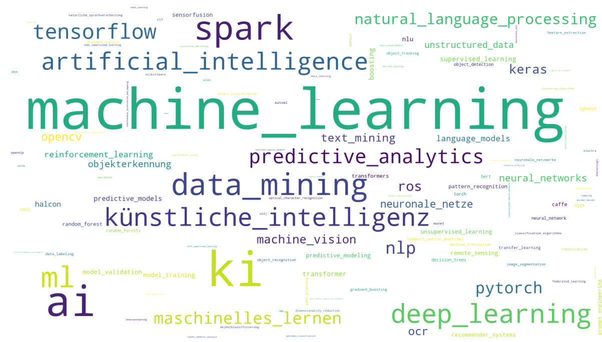

(b) AI Methods

NOTE. —This word cloud comprises keywords that are associated with AI applications and AI methods. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2024; own calculations.

Figure B2 visualizes the resulting domains with separate word clouds, while Figure B3 shows trends in AI skill demand by domain from 2017–2023. Our comprehensive measure indicates a slight increase in AI skill demand between 2017 and 2023, though with notable fluctuations and a downward trend between 2021 and 2023. Panel B3b highlights that this pattern is driven mainly by demand for developer-stage skills ("AI methods"), while demand for application-stage skills has been rather flat (Panel B3c).

Figure B3: Trends in Domains of AI Demand: Methods and Applications (2017–2023)

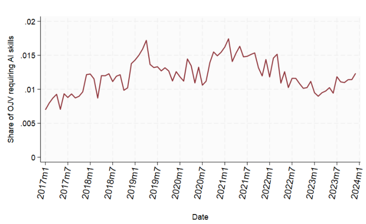

(a) Baseline

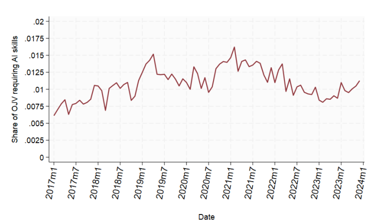

(b) AI Methods

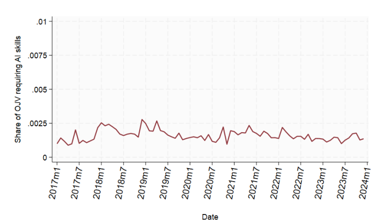

(c) AI Applications

NOTE. —Figure B3 provides an overview of AI trends using the comprehensive baseline taxonomy. Panel B4a is our baseline AI measure, Panel B3b focuses on AI methods, while Panel B3c displays AI applications. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; own calculations.

A2.3 AI measure: External Validity

While our AI taxonomy is grounded in the existing literature, it is more comprehensive due to our enrichment outlined in section A2.1. This enrichment procedure could raise concerns about potential false positives, i.e. erroneously classifying a vacancy as an AI vacancy. Restricting the identification of AI skills analysis to job profile sections mitigates this risk, but we provide further evidence on the validity of our measure.

Figure B4: Trends in AI Demand: Alternative AI Measures (2017–2023)

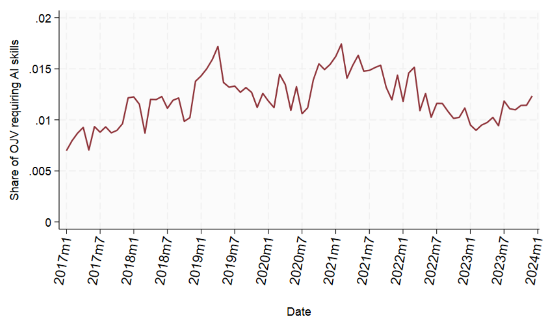

(a) Baseline

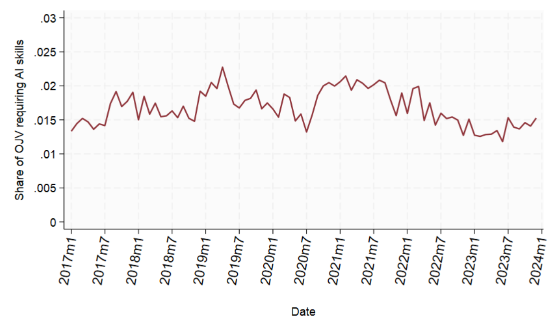

(b) Büchel et al. (2021)

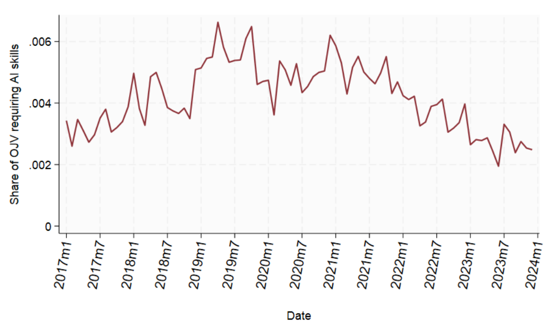

(c) Acemoglu et al. (2022)

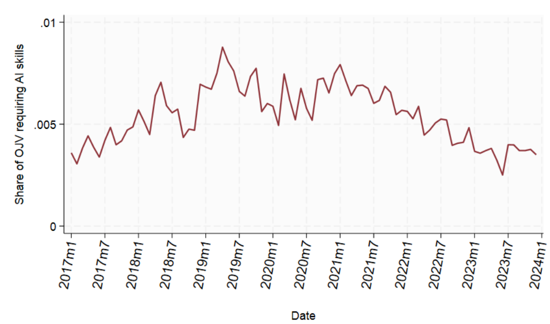

(d) Lightcast

NOTE. —Panel B4a displays AI skill demand over time using our baseline taxonomy. Panel B4b adopts the Büchel et al. (2021) taxonomy, Panel B4c uses the taxonomy from Acemoglu, Autor, Hazell, and Restrepo (2022), and Panel B4d applies the Lightcast taxonomy. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; own calculations.

To validate our AI measure, we construct alternative definitions using the exact taxonomies from (i) Lightcast, (ii) Acemoglu, Autor, Hazell, and Restrepo (2022), and (iii) Büchel, Demary, Goecke, Kohlisch, Koppel, Mertens, Rusche, Scheufen, and Wendt (2021).

Figure B4 compares trends from these measures with our own, revealing two key insights. First, our measure shows more stable demand for AI skills than the AAHR and Lightcast taxonomies, both of which display a downward trend since 2017. Second, our baseline measure more closely tracks stylized facts based on the Büchel et al. (2021) taxonomy, whose inclusion of German keywords makes it particularly relevant for our context.

We also compare our AI measure to the Indeed Hiring Lab's AI Tracker (Indeed Hiring Lab 2024), which provides AI vacancy shares in Germany from 2019–2023.[5] Figure B5 shows that both indicators capture similar dynamics: a slight pre-pandemic decline, a sharp increase in mid-2020, and a steady decline starting in late 2021.

Figure B5: External Validity on AI Vacancies Over Time: Finbot vs Indeed (2017–2023)

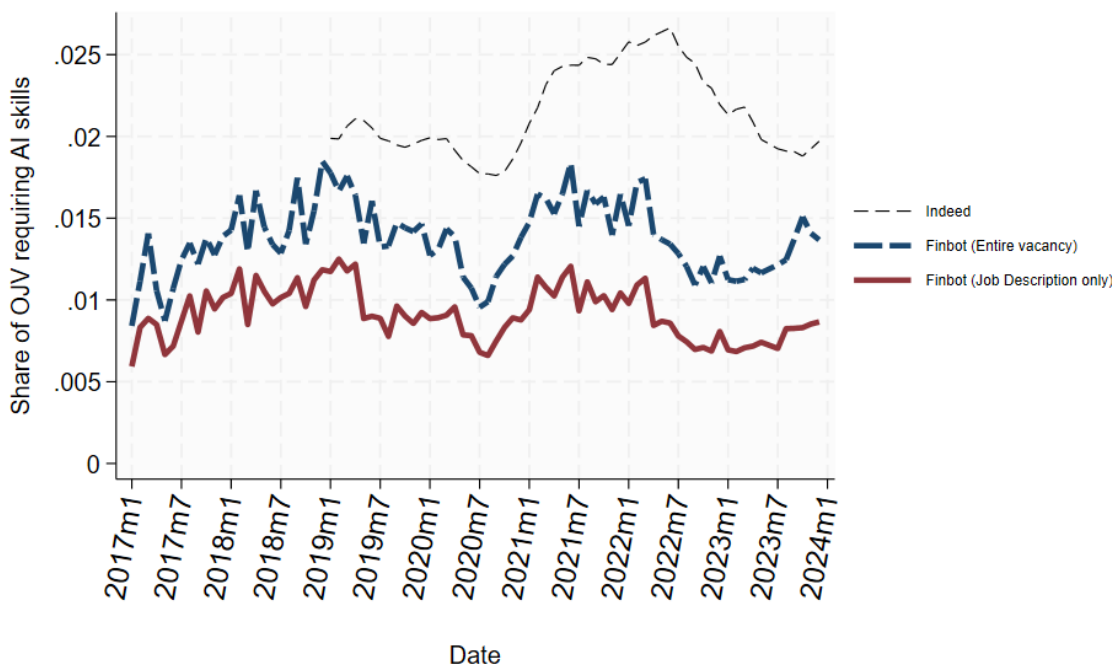

NOTE. —This figure compares the share of AI vacancies of our own OJV data with similar data from Indeed Hiring Lab's AI Tracker. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023, Indeed Hiring Lab (2024); own calculations.

[5]The Indeed AI tracker is available at: `https://github.com/hiring-lab/ai-tracker` . We thank Pawel Adrjan from Indeed for pointing us to this free and publicly available data source.

Finally, we also provide external validity of our vacancy-based measure against the widely used Felten, Raj, and Seamans (2021) measure. Figure B6 shows that our vacancy-based measure aligns closely with the occupational AI exposure index from Felten, Raj, and Seamans (2021). Occupations ranked highly in the Felten et al. index —such as occupations in computer science, economics, or business organization —also exhibit the highest AI vacancy shares in our data. The strong correlation ($\rho = 0.69$) suggests that our vacancy-based approach captures similar patterns of AI exposure as the "AI feasibility" embedded in the Felten et al. measure. Taken together, these comparisons reassure us that our baseline AI measure is robust and consistent with external benchmarks.

Figure B6: External Validation of OJV-based AI Measure: Felten (Rank-Normalized)

NOTE. — This figure compares the OJV-based AI exposure measure with the external index from Felten, Raj, and Seamans (2021), both normalized to percentile ranks (0—1). The fitted line indicates a positive and strong association between both measures ($\rho = 0.69$). Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

A3 Details on IV validation

Table B3: Balance Tests: Correlation between individual- and occupation-specific Characteristics with AI skill demand

	Panel A: Unit-Level (Individual, 2012)								
	Education	Age	Tenure	Experience	Skill	Foreign	Male	Firm Size	Industry
AI Exposure	0.0507	1.3690	3.2852	-3.3612**	-0.3831	-0.0702	-0.1662	0.2603	-0.5540
	(0.1724)	(3.3444)	(2.9114)	(1.3432)	(0.3031)	(0.1629)	(0.1291)	(0.4213)	(0.6732)
Observations	302,768								
	Panel B: Shift-Level (Occ. (3-digit), 2012)								
	Mean Educ.	Mean Age	Mean Tenure	Mean Exp.	Mean Skill	Share Foreign	Share Male	Avg. Firm Size Cat.	Avg. Indus. Cat.
AI Exposure	0.0586***	-0.1888	-0.0726	-0.2436	0.1128***	-0.0017	0.0141	0.0533**	0.0763
	(0.0155)	(0.1626)	(0.1186)	(0.1474)	(0.0356)	(0.0039)	(0.0135)	(0.0229)	(0.0620)
Observations	118								

NOTE. — The respective dependent variable is excluded from the set of controls when used as the outcome. Standard errors in parentheses, clustered at the occupation-region level where applicable. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

Table B4: IV Regressions of Past Earnings and Employment on AI Skill Demand (Pre-Trends): Placebo Tests

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Panel A: Days Employed (Past)								
AI Exposure	-0.193**	-0.002	-0.002	-0.003	0.099	0.083	0.082	0.050	0.050
	(0.090)	(0.076)	(0.076)	(0.074)	(0.068)	(0.067)	(0.067)	(0.062)	(0.062)
	Panel B: Log Annual Earnings (Past)								
AI Exposure	0.005***	0.003***	0.003***	0.003***	0.003***	0.001*	0.001*	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,648,537								
Controls: Socio		✓	✓	✓	✓	✓	✓	✓	✓
Controls: Work			✓	✓	✓	✓	✓	✓	✓
Controls: Firm				✓	✓	✓	✓	✓	✓
State × Year FE					✓		✓		✓
Occ. (1-digit) ×Year FE						✓	✓		
Occ. (2-digit) × Year FE								✓	✓

NOTE. — IV regressions of outcomes from the past (2012–2016) on AI exposure (2017–2021). All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Standard errors in parentheses, clustered at the occupation-region level. Significance Levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2021; SIAB Version 7523 v1; own calculations.

A4 Details on Robustness Checks

In this section we provide a detailed discussion of our robustness checks to ensure validity of our main results. In section A4.1 we repeat our baseline approach using alternative AI measures. Section A4.2 explores the role of mobility patterns, which could imply selection into labor markets as a driving force of our baseline results. In section A4.3 we test the sensitivity of our estimates to alternative model specifications, including non-linear specifications. In section A4.4 we check for potential biases due to wage censoring and sample selection. Lastly, in section A4.5 we provide results based on a conventional shift-share instrument.

A4.1 Alternative Definitions of AI Demand

One concern with our baseline model is mismeasurement of AI skill demand. We address this issue by drawing on our domain-specific measures introduced in section A2.2 (Internal Validity) and AI taxonomies from the literature outlined in section A2.3 (External Validity).

I. Internal Validity

Different AI domains may have distinct labor market implications (Aghion, Bunel, Jaravel, Mikaelsen, Roulet, and Søgaard 2025). To test if these differences affect our baseline estimates, we re-estimate our main models using separate measures for (i) AI methods (developer-stage) and (ii) AI applications (industry-specific use cases). Table C1 reports the results.

Table C1 shows that our earnings results are primarily driven by AI methods. Here, a doubling in the share of AI vacancies raises annual earnings by 0.67%, compared to 0.71% in the baseline. In contrast, earnings gains associated with AI applications are very minor (0.09%). Employment responses follow the same pattern. Both the baseline (–0.15 days) and domain-specific results (–0.11 for methods; –0.09 for applications) point to very small displacement effects at most, though the magnitudes are weaker once domains are separated.

Overall, these results indicate that our modest positive baseline effects on earnings are driven primarily by demand for developer-stage AI skills, while industry-specific applications play a subordinate role. At the same time, employment effects are consistently minor across both domains, reinforcing our main conclusion that AI exposure has negligible implications for employment stability on average.

Table C1: IV Regressions of Days Employed and Earnings on AI Skill Demand: Alternative AI Measures (Internal Validity)

	Days Employed			Log Annual Earnings		
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Methods	Applications	Baseline	Methods	Applications
AI Skill Demand	-0.1481***	-0.1065***	-0.0944***	0.0071***	0.0067***	0.0009***
	(0.0363)	(0.0328)	(0.0231)	(0.0002)	(0.0002)	(0.0001)
Observations		2,385,614			2,385,614	

NOTE. — All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

II. External Validity

To assess whether our findings depend on the choice of our own taxonomy, we replicate our baseline analysis using alternative AI definitions from the literature (see Section A2.3). Specifically, we compare our taxonomy to (i) Lightcast, (ii) Büchel et al. (2021), which includes German-language AI keywords, and (iii) Acemoglu, Autor, Hazell, and Restrepo (2022), a selective adaptation of Lightcast.

Table C2 provides two key insights. First, all measures yield negative employment estimates, though magnitudes differ. Our taxonomy (–0.15 days), Büchel (–0.16), and Lightcast (–0.06) are statistically significant, while AAHR (–0.01) is not. Second and similarly, earnings results are uniformly positive but vary in size. The estimates range from 0.46% (AAHR) to 0.71% (our taxonomy). These differences reflect the scope of the underlying keyword lists. For example, the Lightcast and Acemoglu, Autor, Hazell, and Restrepo (2022) taxonomy do not include German keywords, while Büchel et al. (2021) does include German keywords. However, the latter taxonomy covers a broader set of adjacent terms such as robotics, big data, and business intelligence which we decided not to include (because not deemed AI skills in a narrow sense) —instead emphasizing recent algorithmic and model-specific developments, including a more extensive set of AI tools and the emergence of Generative AI.

Table C2: IV Regressions of Days Employed and Earnings on AI Skill Demand: Alternative AI Measures (External Validity)

	Days Employed				Log Annual Earnings			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Lightcast	Büchel et al. 2021	AAHR	Baseline	Lightcast	Büchel et al. 2021	AAHR
AI Skill Demand	-0.1481***	-0.0567**	-0.1570***	-0.0114	0.0071***	0.0055***	0.0062***	0.0046***
	(0.0363)	(0.0230)	(0.0469)	(0.0206)	(0.0002)	(0.0003)	(0.0002)	(0.0002)
Observations	2,385,614				2,385,614			

NOTE. — All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

Taken together, the results show that our findings are robust across alternative taxonomies, while also underscoring the importance of careful definitional choices when measuring AI skill demand from vacancy data (in terms of quantitative differences).

A4.2 Mobility and Selection into Labor Markets

Through the lens of the AAHR model, positive responses to AI demand are interpreted as "on-the-job" gains, resulting from productivity-enhancing features of AI in workers' current jobs. Yet, compositional effects could distort this interpretation if gains are instead driven by workers moving across regions or occupations, i.e. "job-to-job" gains. To assess this possibility, we vary the FE structure of our baseline model with three differentially restricted specifications: (i) only 3-digit occupation FE ("Region mover"), capturing workers who switch regions but stay in the same occupation; (ii) only CZ FE ("Occupation mover"), capturing workers who switch occupations but stay in the same region; and (iii) no occupation nor CZ FE ("Region + Occupation mover"), capturing workers who change both.

Table C3 shows that earnings estimates remain positive across all specifications, though smaller when we restrict the sample to "Region + Occupation movers" (+ 0.3%) compared to the baseline results (+ 0.7%). For employment, we find negative estimates in the baseline analysis (–0.15 days) and for "Region movers" (–0.15), but small positive estimates for

"Occupation mover" and "Region + Occupation movers" (+0.07 days each). These patterns suggest that the modest earnings gains we observe are not primarily driven by job-to-job mobility. While we do find sign changes for employment once we omit occupation FE, we caution against over-interpreting these results, as the uniformly small magnitudes indicate that AI exposure has negligible implications for employment stability in general.

Combined, this evidence does not support mobility as a central channel of AI-induced labor market responses. Instead, our findings remain consistent with "on-the-job" productivity gains, in line with Gathmann, Grimm, and Winkler (2024), who also find little evidence that AI exposure triggers occupational mobility or generates earnings gains for job movers.

Table C3: IV Regressions of Days Employed and Earnings on AI Skill Demand: Mobility Restrictions and Labor Market Selection

	(1) Baseline	(2) Region Mover	(3) Occupation Mover	(4) Region + Occ. Mover
	Days Employed			
AI Skill Demand	-0.1481***	-0.1476***	0.0724***	0.0735***
	(0.0363)	(0.0363)	(0.0262)	(0.0262)
	Log Annual Earnings			
AI Skill Demand	0.0071***	0.0071***	0.0030***	0.0030***
	(0.0002)	(0.0002)	(0.0001)	(0.0001)
Commuting Zone FE	✓		✓	
Occ. (3-digit) FE	✓	✓		
Observations	2,385,614			

NOTE. — All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

A4.3 Model Specification

We next test the robustness of our results to alternative model specifications. Table C4 reports employment and earnings estimates and shows that our baseline results remain robust to alternative scaling choices of our AI skill measure. We compare our baseline normalization (interpretable as a 100% increase in AI demand) with two alternatives: a share-based measure (1 pp. increase) and an intensity-based measure (average number of AI skills per posting). Both produce qualitatively similar results. The share-based measure implies that a 1 pp. increase in AI exposure raises earnings by 0.6% (baseline: 0.7%), while employment results remain small and insignificant. The intensity-based measure yields negligible effects, suggesting that AI exposure is driven more by the binary presence of AI skills than by the breadth of skills listed. We also rescale our AI measure using the IHS transformation, as in Acemoglu et al. (2022). This transformation implies much stronger earnings gains (4.2%), but we remain cautious in interpreting these magnitudes given known issues with IHS variables and their sensitivity to zeros.[6]

Next, we test directly for non-linearities by adding a squared term to our baseline specification. For employment, the linear coefficient is negative (–0.28) while the squared term is positive (+0.005). This pattern implies that workers in low-AI exposure labor markets, who are more likely to be lesser-skilled, face small displacement risks, whereas those in high-AI exposure markets, with a greater share of skilled workers, experience more stable employment. For earnings, we find the reverse pattern: the linear coefficient is positive (+1.0%) while the squared term is negative (–0.01%). This pattern suggests that earnings gains are strongest in low- to medium-exposure labor markets but flatten out in highly exposed ones. These dynamics reinforce our heterogeneity results, showing that the impact of AI on employment and earnings depends on skill composition and task structure of affected labor markets.

Finally, as an additional check, we replicate our baseline analysis using an "extended" AI measure based on the full vacancy text rather than the job profile alone. This broader measure yields somewhat larger earnings gains (1%), but no significant impact on employment. While informative, we interpret these results cautiously, since they are more likely to capture noise from non-job-specific mentions of AI in online job ads.

Overall, these robustness checks suggest that our main results are not sensitive to alter-

[6] IHS transformations approximate log functions for large values and handle zeros, but recent work highlights their sensitivity to distributional features (Chen and Chan 2024; Mullahy and Norton 2024). We therefore prefer our normalized baseline measure, which is both conservative and more intuitive to interpret.

native model specifications. Across scaling approaches and functional forms, AI exposure has modest implications on earnings and negligible ones for employment.

Table C4: IV Regressions of Days Employed and Earnings on AI Skill Demand: Alternative AI Skill Demand Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Days Employed					Log Annual Earnings				
AI: Share-Based (1 pp.)	-0.1179					0.0061***				
	(0.1089)					(0.0006)				
AI: Base (Linear)		-0.2763***					0.0100***			
		(0.0524)					(0.0004)			
AI: Base (Squared)		0.0045***					-0.0001***			
		(0.0013)					(0.0000)			
AI: IHS-Transformed			-0.1918					0.0416***		
			(0.2745)					(0.0014)		
AI: Intensity				-0.0367					0.0010***	
				(0.0319)					(0.0002)	
AI: Entire vacancy					0.0998					0.0100***
					(0.0730)					(0.0004)
F-statistic (1st stage)	872.1	9.7	10,672.3	624.9	2,474.8	871.8	9.7	10,664.5	624.8	2,473.4
Observations	2,385,502	2,385,614	2,385,502	2,385,502	2,385,502	2,385,502	2,385,614	2,385,502	2,385,502	2,385,502

NOTE. — All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Alternative measures of AI exposure: share-based (1 pp increase), non-linear, IHS-transformed, intensity-based, and share-based using the entire vacancy text. Standard errors in parentheses, clustered at the CZ level. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

A4.4 Sample selection & Wage censoring

Next, we test whether our results are sensitive to sample restrictions in the SIAB and OJV data or to the treatment of top-coded wages (Tables C5 and C6). These restrictions reduce statistical power but increase consistency.

Table C5: IV Regressions of Days Employed and Earnings on AI Skill Demand: Alternative Sample Restrictions

	(1)	(2)	(3)	(4)	(5)
	Days Employed				
AI Demand	-0.1481***	0.1940***	-0.1449***	0.1969***	-0.2026***
	(0.0363)	(0.0339)	(0.0362)	(0.0339)	(0.0426)
	Log Annual Earnings				
AI Demand	0.0071***	0.0081***	0.0072***	0.0081***	0.0104***
	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0002)
Baseline	✓				
Balanced Worker Sample		✓		✓	
Balanced CZ Sample			✓	✓	
Winsorized AI Skill Demand (Top 1%)					✓
AI Share Mean	0.006				
Observations	2,385,614	1,914,363	2,332,232	1,907,727	2,385,614

NOTE. —All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

In the first step we restrict the sample along four dimensions: (i) workers present in the administrative data throughout 2017–23, (ii) workers observed for more than four years, (iii) CZ–occupation cells with at least three postings each year, and (iv) a combination of (i) and (iii). A fifth specification uses a winsorized AI exposure measure at the 99th percentile to test if results are driven by a few highly exposed cells. Across all cases, earnings estimates remain modestly positive. For employment, more balanced worker samples unsurprisingly produces small positive estimates as continuously observed workers are more likely to be steadily employed. The winsorized measure yields nearly identical employment results but somewhat larger earnings gains (+0.1%), implying our baseline (+0.7%) is a lower bound.

In the second step we test the sensitivity of our results to wage censoring. In the main analysis, we use imputed wages following Fitzenberger, Osikominu, and Völter (2006) for the calculation of annual earnings. Re-estimating the models with top-coded wages flips the baseline earnings gains (+0.7%) to earnings losses (–0.2%). These contrasting results underline that earnings gains are concentrated among high-wage workers whose incomes would otherwise be censored.

Table C6: IV Regressions of Earnings and Employment on AI Skill Demand: Imputed and Censored Wages and Earnings

	(1)	(2)	(3)	(4)
	Log Wages		Log Annual Earnings	
	Imputed	Censored	Imputed	Censored
AI Demand	0.0042***	-0.0020***	0.0071***	-0.0019***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Observations	2,385,614	2,119,374	2,385,614	2,119,374

NOTE. —All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data) 2017–2023; SIAB Version 7523 v1; own calculations.

A4.5 Alternative IV measures

As a final robustness check, we complement our baseline LOOM IV with (i) an extended LOOM IV in which we account for possible regional spillovers and (ii) a more conventional shift-share instrument.

First, we assess the potential influence of spatial spillovers. We construct an alternative version of our LOOM instrument that excludes not only the worker's own commuting zone but also all adjacent CZs when computing the leave-one-out mean. This modification alleviates concerns that our baseline AI exposure may be contaminated by spillovers due to broader technological trends. As shown in Table C7, results remain virtually unchanged compared to our baseline specification, implying that broader spillover effects are unlikely to confound our main identification strategy.

Table C7: IV Baseline Regressions of Days Employed and Earnings on AI Skill Demand: Spillover-Resistant LOOM IV

	(1)	(2)	(3)	(4)	(5)
	Days Employed				
AI Demand	-0.2741***	-0.1576***	-0.1587***	-0.1575***	0.0760*
	(0.0353)	(0.0353)	(0.0353)	(0.0353)	(0.0444)
F-statistic (1st stage)	4,059.1	3,689.1	3,689.0	3,689.2	1,578.8
	Log Annual Earnings				
AI Demand	0.0081***	0.0071***	0.0071***	0.0071***	-0.0000
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0002)
F-statistic (1st stage)	3,694.2	3,688.0	3,688.0	3,688.2	1,578.8
Controls: Socio		✓	✓	✓	✓
Controls: Work			✓	✓	✓
Controls: Firm				✓	✓
Occ. (2-digit) × Year FE					✓
AI Share Mean	0.006				
Observations	2,385,616				

NOTE. —All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB v7523; own calculations.

Finally, we construct a conventional shift-share instrument. We follow standard practice (Borusyak, Hull, and Jaravel 2025), using occupation-CZ employment shares from 2012 as base shares and interact them with national shocks to AI demand at the occupation level (while still excluding AI demand in worker's own CZ). The identifying variation thus comes from shifts in national AI demand, apportioned across CZs according to their occupational structure in 2012.

Table C8: IV Baseline Regressions of Days Employed and Earnings on AI Skill Demand: Shift-Share

	(1)	(2)	(3)	(4)	(5)
	Days Employed				
AI Demand	-0.1088***	-0.1069***	-0.1083***	-0.1074***	0.0816*
	(0.0355)	(0.0361)	(0.0361)	(0.0361)	(0.0470)
F-statistic (1st stage)	3,861.2	3,620.5	3,620.5	3,620.4	1,419.1
	Log Annual Earnings				
AI Demand	0.0079***	0.0066***	0.0066***	0.0066***	0.0001
	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
F-statistic (1st stage)	3,627.8	3,619.2	3,619.2	3,619.2	1,419.0
Controls: Socio		✓	✓	✓	✓
Controls: Work			✓	✓	✓
Controls: Firm				✓	✓
Occ. (2-digit) × Year FE					✓
AI Share Mean	0.006				
Observations	2,385,614				

NOTE. —All regressions are based on equation (3) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB v7523; own calculations.

Table C8 shows the results closely mirror our main analysis. For earnings, coefficients are positive in simpler specifications without cross FE (columns 1–4). In response to a doubling in the share of AI vacancies, earnings range from 0.66% to 0.79%. As in our main analysis, these positive estimates vanish once we include occupation-by-year FE (column 5). For employment, estimates are modestly negative across simpler specifications, implying that a doubling in the share of AI vacancies reduces annual employment by about a tenth of a day. Once occupation-by-year FE are included, however, these estimates turn weakly positive (+0.08 days), echoing our main results.

Overall, the shift-share design confirms that AI skill demand has negligible implications for employment stability and only modest earnings effects for the average worker. The similarity of estimates across LOOM and shift-share designs thus suggests that our findings are not driven by instrument choice but reflect genuine, though limited, labor market impacts of AI demand.

A5 Details on Extraction of Skills from BERUFENET

This section outlines how we construct our dataset on non-AI skill requirements, which is central to the analysis on reinstatement effects of new AI technologies (Acemoglu, Autor, Hazell, and Restrepo 2022). To this end, we transform raw information from the job portal BERUFENET into a data structure that is suitable for econometric analysis. BERUFENET, a database analogous to O*NET and ESCO, provides a rich source of job requirements, encompassing tasks, skills, certificates, and technologies.

In a first step, we extract about 8,700 raw job requirements directly from BERUFENET. These entries, however, blend concepts such as skills, tasks, certificates, and technologies. Guided by our conceptual framework (section 2), we treat skills, tasks, and certificates as worker qualifications, while technologies are tools (e.g., Tableau) applied in production. We view this focus on skills and tasks more in line with the conceptual framework laid out in Acemoglu, Autor, Hazell, and Restrepo (2022), according to which new technologies expand or shrink the job-specific task space. To ensure consistency, we use ChatGPT-4o to separate the two categories. In the following, we consider skills, tasks, and certificates interchangeable —as all these segments constitute the job description —and refer to these simply as "skills".

Second, we enrich each skill with contextual descriptions and example occupations, similar to Arntz, Böhm, Graetz, Gregory, Lehmer, and Lipowski (2025). This step helps disambiguate terms to enhance the quality of the ensuing classification. For example, the German job task “betreuen” can be used in the context of serving professional clients (then NR interactive) or in the context of, say, caretaking of the elderly (then NR manual). Adding job descriptions and example occupations helps distinguishing between related, but (in our view) different skill requirements.

Third, we classify all skills into five established categories (see **Storm.2022**; **Storm.2022b** for a discussion): non-routine (NR) analytic, NR interactive, routine (R) cognitive, R manual, and NR manual. To accommodate this classification, we prompt ChatGPT to categorize skills, building upon explicit decision rules outlined in Dengler, Matthes, and Paulus (2014). Subsequently, we iteratively validate this initial classification together with research assistants to address ambiguities and make human decisions on edge cases. About 13% of initial classifications are manually reclassified, often cases involving IT-related activities with varying degrees of repetitiveness and activities involving lots of personal interactions. For example, repetitive activities such as IT-admin are considered routine cognitive, while more complex activities are considered NR analytic. Similarly, activities with many personal interactions in the business context are typically classified as NR interactive, while activities with many personal interactions in the hospitality and health sector are often classified as NR manual.

Fourth, we perform various linguistic adjustments to ensure that our set of keywords captures the linguistic diversity seen in job postings. A key challenge is that BERUFENET was never created to accommodate data preparation suitable for the analysis of job vacancies or related text data. Instead, BERUFENET is aimed at being an informative source for job seekers, commonly displayed by formal, noun-based descriptions (e.g., ”development” of something), while job postings often use verb forms (e.g., applicants must ”develop” something). To bridge this semantic gap, we systematically include both, noun and verb forms in our keyword lists and we address compound skills that BERUFENET often combines by using hyphens (e.g., Investment-/Financial advice). Firms do not always use this compound structure, often instead opting for more specific requirements (financial advice and/or investment advice). Combining such compound skill requirements with separate entries therefore enhances the precision of our search algorithm.

Table D1: Skill Classification and Example Occupations

Skill	**Group**	**Skill Description**	**Occupation**
Data Analysis	NRA	Analyzing data to extract meaningful insights for decision-making.	Data Analyst
Commercial Law	NRA	Applying laws and regulations related to commerce.	Commercial Lawyer
Financial Advisory	NRI	Providing investment and financial planning advice.	Financial Advisor
Private Tutoring	NRI	Teaching individuals in a one-on-one setting.	Tutor
Carpentry	NRM	Crafting wood structures and components for buildings.	Carpenter
International Cuisine	NRM	Cooking international cuisine.	Chef
Hardware Installation	RC	Setting up and configuring hardware systems.	Hardware Installer
HR Administration	RC	Managing administrative HR tasks.	HR Administrator
Sheet Metal Fabrication	RM	Fabricating metal sheets into various shapes and products.	Metal Fabricator
Warehouse Work	RM	Performing tasks related to warehouse operations.	Warehouse Worker

NOTE. — This table maps selected skills to their respective task categories used in our analysis: Non-Routine Analytic (NRA), Non-Routine Interactive (NRI), Non-Routine Manual (NRM), Routine Cognitive (RC), and Routine Manual (RM). The examples serve illustrative purposes and reflect the task content of occupations most frequently associated with each skill.

Our final dataset comprises about 9,500 unique keywords, covering both the breadth of conceptually distinct skills and the depth of linguistic variants. Figure D1 shows corresponding word clouds and provides intuitive insights. For example, *management* and *development* are important activities within the skill group NR analytic, while activities associated with *consultation* and *sales* are the most important activities within the task group NR interactive. Regarding the routine skill groups, technical competencies such as *mechatronics* or basic language skills such as *German* are important activities within routine cognitive and *logistics* and *production* related activities are important within routine manual. Lastly, *maintenance* and *caretaking* related activities are important activities within NR manual.

Table D2 further provides an overview of the Top 5 occupations by task intensity.

Table D2: Top 5 Occupations by Task Intensity (per Task Group)

Rank	Occupation *(3-digit KLdB 2010)*	Task Intensity	Job ads (in k)
	Panel A: Top 5 occupations – Analytic (NRA)		
1	Product and industrial design *(931)*	75.05%	21.9
2	Technical media design *(232)*	71.67%	52.0
3	Software development and programming *(434)*	71.39%	442.0
4	Laboratory occupations in medicine *(812)*	65.24%	29.9
5	Teachers/researchers at universities *(843)*	64.04%	54.1
	Panel B: Top 5 occupations – Interactive (NRI)		
1	Sales (retail) books, antiques, music *(625)*	61.66%	0.7
2	Sales (retail) foodstuffs *(623)*	55.89%	18.0
3	Sales (retail) durables (e.g. clothing, cars) *(622)*	52.93%	20.9
4	Sales (retail) (general) *(621)*	51.32%	353.8
5	Tourism and sports (fitness) *(631)*	50.65%	7.8
	Panel C: Top 5 occupations – Routine Cognitive (RC)		
1	Inspect. & maint. of traffic infrast. *(512)*	44.06%	3.1
2	Accounting, controlling, auditing *(722)*	41.17%	366.1
3	Office clerks and secretaries *(714)*	38.40%	143.1
4	Hotels *(632)*	36.89%	44.4
5	Public administration *(732)*	36.84%	22.1
	Panel D: Top 5 occupations – Routine Manual (RM)		
1	Construction and transport equipment ops. *(525)*	50.14%	23.8
2	Metal construction and welding *(244)*	46.58%	42.1
3	Wood-working and processing *(223)*	44.37%	15.5
4	Interior construction, carpentry, glazing *(333)*	40.05%	8.2
5	Warehousing, logistics, postal delivery *(513)*	38.41%	178.2
	Panel E: Top 5 occupations – Non-Routine Manual (NRM)		
1	Cleaning services *(541)*	52.97%	80.4
2	Geriatric care *(821)*	40.52%	52.0
3	Floor laying *(331)*	37.58%	5.2
4	Housekeeping and consumer counselling *(832)*	32.89%	19.3
5	Technical automotive/aero/shipbuilding *(252)*	29.95%	66.2

NOTE. — This table lists the Top 5 occupations (3-digit KLdB 2010 level) by task intensity for each of the five task categories: Non-Routine Analytic (NRA), Non-Routine Interactive (NRI), Non-Routine Manual (NRM), Routine Cognitive (RC), and Routine Manual (RM). Job ads refer to the total number of job ads for each occupation between 2017 and 2023, rounded to the nearest hundred. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; own calculations.

Figure D1: Word Clouds of (non-AI) Skills by Task Category

Non-Routine Analytic　　Non-Routine Interactive

Non-Routine Manual　　Routine Cognitive　　Routine Manual

NOTE. — This figure shows word clouds of most frequently stated skill and task requirements in online job ads. Larger terms occur more frequently. Source: Own calculations based on Palturai GmbH/Finbot AG (OJV data), 2017–2024.

To validate our classification, we compare it with the IAB Occupational Panel (Grienberger, Janser, and Lehmer 2023). This data also builds on BERUFENET and provides information on the occupation-level task structure. Because of the same underlying data source, we view this comparison as informative on the validity of our classification procedure. To that end, we compute task intensity measures T_{io} as the share of each task type in occupation o (i.e. the same definition of task structures provided in the IAB Occupational Panel):

$$T_{jo} = \frac{\text{Number of tasks j demanded in occupation o}}{\text{Total number of tasks demanded in occupation o}} \tag{9}$$

where $j = 1, ..., 5$ represents the five tasks. This definition implies (i) $T_{ijlmt} \in [0, 1]$ $\forall j$ and (ii) $\sum_J T_{ijlmt} = 1$, thus the relative importance of each task j in occupation o.

First, we correlate our OJV-based task intensities with those from the IAB panel (Table D3). We find high correlations across all categories, especially for NR analytic (0.75) and NR interactive (0.75). Our measures for Routine manual (0.69), NR manual (0.59), and Routine cognitive (0.50) also show substantial correlation. Note that perfect correlation is not expected since our data capture within-occupation heterogeneity, while the IAB panel reflects occupation averages only.

Table D3: Correlation between vacancy-based task composition & IAB Occupational Panel: Occupational Task Structure

	OJV data				
	NRA	NRI	RC	RM	NRM
IAB Panel					
NRA	0.75				
NRI		0.75			
RC			0.50		
RM				0.69	
NRM					0.59

NOTE. — This table shows correlations between the occupational task composition found in online job ads and the IAB Occupational Panel, aggregated at the 3-digit occupation level.

Table D4: Correlation between vacancy-based task composition & IAB Occupational Panel: & IAB Occupational Panel: Substitution Potential

	OJV data				
	NRA	NRI	RC	RM	NRM
IAB Panel					
NRA	-0.36				
NRI		-0.33			
RC			0.40		
RM				0.45	
NRM					0.15

NOTE. — This table reports correlations between the occupational task composition found in online job ads and the IAB Occupational Panel's measure of substitution potential, matched at the 3-digit occupation level. Negative values (e.g., NRA, NRI) suggest that occupations with higher intensity in these tasks are less likely to be substituted.

Second, we compare our measures to indicators of substitution potential (SP) (Dengler and Matthes 2018). The SP captures how easily tasks within an occupation can be automated by new technologies. Positive correlations indicate higher substitutability, while negative correlations suggest complementarities with technological adoption. As shown in Table D4, both routine categories correlate positively with SP (0.40 – 0.45), consistent with their higher automation risk. In contrast, NR analytic (–0.36) and NR interactive (–0.33) correlate negatively, reflecting complementarities with technology. NR manual shows a weak positive correlation (0.15), implying limited substitutability but also limited complementarities.

Taken together, these validation exercises confirm the robustness of our task classification. Our OJV-based measures align closely with the IAB Occupational Panel and yield intuitive patterns of substitutability and complementarity with technologies.

A6 Heterogeneity: Detailed Results

Table E1: IV Regressions of Earnings and Employment on AI Skill Demand: By Skill Groups

	(1) Days Employed	(2) Log Annual Earnings
	Panel A: Baseline AI Measure	
Unskilled	-0.1870**	-0.0033***
	(0.0934)	(0.0005)
Skilled	0.1093	-0.0029***
	(0.0718)	(0.0004)
Specialists	0.6452***	0.0013
	(0.1935)	(0.0010)
Experts	0.8654***	0.0065***
	(0.2496)	(0.0012)
	Panel B: AI Methods	
Unskilled	-0.2378***	-0.0032***
	(0.1064)	(0.0006)
Skilled	0.0605	-0.0030***
	(0.0589)	(0.0003)
Specialists	0.4857***	-0.0009
	(0.1678)	(0.0008)
Experts	0.6558***	0.0030***
	(0.2071)	(0.0010)
	Panel C: AI Applications	
Unskilled	-0.1791*	-0.0024***
	(0.0970)	(0.0006)
Skilled	-0.1536	-0.0035***
	(0.1381)	(0.0009)
Specialists	-0.0243	0.0069***
	(0.3119)	(0.0021)
Experts	-0.2126	0.0175***
	(0.5161)	(0.0035)
Observations	2,385,502	2,385,502

NOTE. —All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. skill level). Earnings regressions also include year FE. Different Panels show different AI measure domains. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017-–2023; SIAB Version 7523 v1; own calculations.

Table E2: IV Regressions of Days Earnings and Employment on AI Skill Demand: By Education

	(1) Days Employed	(2) Log Annual Earnings
No Degree	-0.3906***	-0.0087***
	(0.1160)	(0.0007)
Vocational Degree	0.0939	-0.0012***
	(0.0822)	(0.0004)
College	0.8209***	0.0043***
	(0.2284)	(0.0011)
Observations	2,385,502	2,385,502

NOTE. —All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. education). Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Table E3: IV Regressions of Earnings and Employment on AI Skill Demand: By Age Groups

	(1) Days Employed	(2) Log Annual Earnings
Young (18–29)	0.7231***	0.0021***
	(0.1278)	(0.0008)
Mid-aged (30–49)	0.7438***	0.0089***
	(0.1365)	(0.0009)
Old (50–65)	0.3441***	-0.0005
	(0.0917)	(0.0005)
Observations	2,385,502	2,385,502

NOTE. —Age groups are defined as young ($18-29$ years old), mid-aged ($30-49$ years), and old ($50-65$ years). All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. age info). Earnings regressions also include year FE. Robust standard errors in parentheses. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

Table E4: IV Regressions of Earnings and Employment on AI Skill Demand: By Labor Market Experience

	(1) Days Employed	(2) Log Annual Earnings
	Panel A: Baseline AI Measure	
Juniors (1–4 yrs)	0.8721***	-0.0009
	(0.1542)	(0.0007)
Experienced (5–14 yrs)	0.5440***	0.0034***
	(0.1398)	(0.0006)
Seniors (15+ yrs)	0.2512**	-0.0004
	(0.0979)	(0.0004)
	Panel B: AI Methods	
Juniors (1–4 yrs)	0.7332***	-0.0013**
	(0.1307)	(0.0006)
Experienced (5–14 yrs)	0.4150***	0.0022***
	(0.1147)	(0.0006)
Seniors (15+ yrs)	0.1797**	-0.0012***
	(0.0792)	(0.0004)
	Panel C: AI Applications	
Juniors (1–4 yrs)	0.0614	-0.0065***
	(0.2089)	(0.0012)
Experienced (5–14 yrs)	-0.1631	-0.0022
	(0.2428)	(0.0014)
Seniors (15+ yrs)	-0.2940	-0.0051***
	(0.2084)	(0.0012)
Observations	2,385,502	2,385,502

NOTE. —All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. LM experience). Earnings regressions also include year FE. Different Panels show different AI measure domains. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Table E5: IV Regressions of Earnings and Employment on AI Skill Demand: By Occupational Task Structure

Panel A: Baseline AI Measure			**Panel B: AI Methods**		
	(1)	(2)		(1)	(2)
	Days Employed	Log Annual Earnings		Days Employed	Log Annual Earnings
Routine Manual	0.3006	-0.0047***	Routine Manual	0.1184	-0.0069***
	(0.1973)	(0.0010)		(0.1751)	(0.0012)
Non-Routine Manual	0.3482***	-0.0032***	Non-Routine Manual	0.1916*	-0.0019***
	(0.1216)	(0.0006)		(0.1023)	(0.0005)
Routine Cognitive	0.0395	0.0007**	Routine Cognitive	-0.0513	-0.0002
	(0.0639)	(0.0003)		(0.0486)	(0.0002)
Non-Routine Interactive	-0.2255**	0.0039***	Non-Routine Interactive	-0.1959**	0.0036***
	(0.0920)	(0.0005)		(0.0788)	(0.0004)
Non-Routine Analytic	-0.1177	0.0061***	Non-Routine Analytic	-0.1165	0.0063***
	(0.1090)	(0.0006)		(0.0947)	(0.0005)
Observations	2,385,502	2,385,502	Observations	2,385,502	2,385,502
Panel C: Applications			**Panel D: Occupational Panel Data**		
	(1)	(2)		(1)	(2)
	Days Employed	Log Annual Earnings		Days Employed	Log Annual Earnings
Routine Manual	1.3935**	0.0144***	Routine Manual	0.1119*	-0.0002
	(0.6054)	(0.0035)		(0.0584)	(0.0003)
Non-Routine Manual	0.5007***	-0.0062***	Non-Routine Manual	-0.2746***	-0.0016***
	(0.1905)	(0.0011)		(0.0871)	(0.0004)
Routine Cognitive	0.1816*	0.0014**	Routine Cognitive	-0.3398***	-0.0006
	(0.0992)	(0.0006)		(0.0942)	(0.0004)
Non-Routine Interactive	-0.2755	0.0016	Non-Routine Interactive	-0.7130**	0.0429***
	(0.2425)	(0.0012)		(0.3530)	(0.0027)
Non-Routine Analytic	0.0743	0.0013	Non-Routine Analytic	0.0966	0.0171***
	(0.2157)	(0.0014)		(0.2528)	(0.0015)
Observations	2,383,027	2,383,027	Observations	2,385,502	2,385,502

NOTE. — All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Different Panels show different AI measure domains. Robust standard errors in parentheses. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2023; SIAB Version 7523 v1; own calculations.

Table E6: IV Regressions of Earnings and Employment on AI Skill Demand: Non-Routine Analytic (NRA) Intensity

	(1) Days Employed	(2) Log Annual Earnings
Non-Routine Analytic Tercile 1	0.0026	-0.0011***
	(0.0692)	(0.0003)
Non-Routine Analytic Tercile 2	-0.0577	0.0018***
	(0.0542)	(0.0003)
Non-Routine Analytic Tercile 3	-0.1758	0.0093**
	(0.1526)	(0.0008)
Observations	2,385,502	2,385,502

NOTE. —All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE. Earnings regressions also include year FE. Estimates are presented separately for occupations in distinct terciles of the NR analytic task distribution, with Tercile 1 being the lowest tercile and Tercile 3 being the highest. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Table E7: IV Regressions of Employment on AI Skill Demand: By Expertise × Occupational Task Intensity

	Days Employed				
	(1) NRA	(2) NRI	(3) RC	(4) RM	(5) NRM
AI Skill Demand × Unskilled	-0.1307	-0.2329**	-0.0988	-0.2846**	-0.1461*
	(0.0970)	(0.1184)	(0.0751)	(0.1232)	(0.0795)
AI Skill Demand × Skilled	-0.2226**	-0.1198**	-0.1554***	-0.1525***	-0.1370***
	(0.1005)	(0.0484)	(0.0467)	(0.0468)	(0.0465)
AI Skill Demand × Specialists	0.2994	-0.0011	-0.0406	-0.0362	-0.0104
	(0.3113)	(0.1258)	(0.1260)	(0.1269)	(0.1255)
AI Skill Demand × Experts	0.2945	0.0371	0.0071	0.0001	0.0298
	(0.6859)	(0.1579)	(0.1594)	(0.1610)	(0.1593)
AI Skill Demand × Unskilled × High	-0.0991	0.1893	-0.3494	0.2021	-0.0782
	(0.1286)	(0.1347)	(0.2263)	(0.1381)	(0.2029)
AI Skill Demand × Skilled × High	0.0866	-0.1581***	0.0083	0.0122	-0.3277*
	(0.0856)	(0.0531)	(0.0585)	(0.1263)	(0.1846)
AI Skill Demand × Specialists × High	-0.3205	-0.0573	0.0970	0.5512	0.0144
	(0.2790)	(0.1380)	(0.1616)	(0.5454)	(0.9715)
AI Skill Demand × Experts × High	-0.2695	-0.0692	-0.0375	0.2791	-0.4386
	(0.6279)	(0.1417)	(0.1512)	(0.3068)	(0.8592)
Observations	2,385,502	2,385,502	2,385,502	2,385,502	2,385,502

NOTE. — All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. skill level and year FE). The task groups considered are Non-Routine Analytical (NRA), Non-Routine Interactive (NRI), Routine Cognitive (RC), Routine Manual (RM), and Non-Routine Manual (NRM). The "High" intensity category corresponds to occupations in the top tercile of the respective task group distribution. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Table E8: IV Regressions of Earnings on AI Skill Demand: By Expertise × Occupational Task Intensity

	Log Annual Earnings				
	(1) NRA	(2) NRI	(3) RC	(4) RM	(5) NRM
AI Skill Demand × Unskilled	-0.0028***	-0.0051***	-0.0022***	-0.0045***	-0.0021***
	(0.0005)	(0.0008)	(0.0004)	(0.0009)	(0.0004)
AI Skill Demand × Skilled	-0.0040***	-0.0001	-0.0007***	-0.0002	-0.0004
	(0.0005)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
AI Skill Demand × Specialists	0.0239***	0.0070***	0.0077***	0.0084***	0.0079***
	(0.0019)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
AI Skill Demand × Experts	0.0404***	0.0139***	0.0136***	0.0145***	0.0140***
	(0.0037)	(0.0010)	(0.0010)	(0.0010)	(0.0010)
AI Skill Demand × Unskilled × High	-0.0013	0.0050***	-0.0051***	0.0020**	-0.0052***
	(0.0009)	(0.0008)	(0.0012)	(0.0009)	(0.0011)
AI Skill Demand × Skilled × High	0.0040***	-0.0028***	0.0003	-0.0042***	-0.0080***
	(0.0005)	(0.0003)	(0.0003)	(0.0007)	(0.0010)
AI Skill Demand × Specialists × High	-0.0160***	0.0090***	-0.0015*	0.0041	0.0259***
	(0.0018)	(0.0008)	(0.0009)	(0.0029)	(0.0053)
AI Skill Demand × Experts × High	-0.0264***	0.0047***	0.0004	-0.0037**	0.0116**
	(0.0034)	(0.0008)	(0.0008)	(0.0017)	(0.0046)
Observations	2,385,502	2,385,502	2,385,502	2,385,502	2,385,502

NOTE. —All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE (excl. skill level). The task groups considered are Non-Routine Analytical (NRA), Non-Routine Interactive (NRI), Routine Cognitive (RC), Routine Manual (RM), and Non-Routine Manual (NRM). The "High" intensity category corresponds to occupations in the top tercile of the respective task group distribution. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own

Table E9: IV Regressions of Net (Non-AI) Skill Change on AI Skill Demand

	(1) All Skills	(2) NRA	(3) NRI	(4) NRM	(5) RC	(6) RM
AI Skill Demand	0.0258**	0.0472***	0.0072**	-0.0149***	-0.0044	-0.0093***
	(0.0112)	(0.0091)	(0.0030)	(0.0021)	(0.0035)	(0.0027)
Observations	65,685	65,685	65,685	65,685	65,685	65,685

NOTE. —The dependent variable in each specification is the net skill change per eq. (6). All regressions are based on equation (7), with task groups: Non-Routine Analytical (NRA), Non-Routine Interactive (NRI), Routine Cognitive (RC), Routine Manual (RM), and Non-Routine Manual (NRM). Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Table E10: IV Regressions of Days Employed and Earnings (RE): By AI Share × Net Changes in Task Demand

	(1) Days Employed	(2) Log Annual Earnings
AI Share	0.4910***	0.0080***
	(0.1722)	(0.0014)
NR Analytic	0.2059***	-0.0017***
	(0.0409)	(0.0003)
NR Interactive	0.1399*	0.0015**
	(0.0782)	(0.0006)
Routine Cognitive	0.1152*	0.0005
	(0.0648)	(0.0005)
Routine Manual	0.0820	-0.0003
	(0.0784)	(0.0005)
NR Manual	0.1918**	-0.0029***
	(0.0814)	(0.0006)
AI Share × NR Analytic	0.0573*	0.0017***
	(0.0334)	(0.0003)
AI Share × NR Interactive	0.1418*	0.0012*
	(0.0809)	(0.0006)
AI Share × Routine Cognitive	0.0609	0.0047***
	(0.0542)	(0.0004)
AI Share × Routine Manual	0.1655	-0.0033***
	(0.1332)	(0.0010)
AI Share × NR Manual	-0.2445	0.0016
	(0.2496)	(0.0019)
AI Share Mean	0.006	0.006
Observations	2,078,928	2,078,928

NOTE. — All regressions are based on equation (8) and include controls described in Section 5, including worker, CZ, and 3-digit-occupation FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017–2022; SIAB Version 7523 v1; own calculations.

Table E11: Earnings Distribution by Wage Decile

Wage Decile	Min Earnings (EUR)	Max Earnings (EUR)
1	71.00	22,065.08
2	22,065.25	27,339.03
3	27,339.04	31,816.34
4	31,816.38	36,087.00
5	36,087.15	40,450.32
6	40,450.32	45,367.53
7	45,367.78	51,789.00
8	51,789.06	61,239.28
9	61,239.30	75,064.27
10	75,064.48	978,684.00

NOTE. —Workers are sorted into deciles based on their observed annual earnings in the main job. Each row shows the minimum and maximum earnings within each bin. Own calculations based on SIAB v7523, 2017–2022.

Table E12: IV Regressions of Earnings and Employment on AI Skill Demand: By Income Bins

	(1) Days Employed	(2) Log Annual Earnings
Income Decile 1	-7.7811***	-0.0387***
	(0.2085)	(0.0011)
Income Decile 2	-3.7459***	-0.0188***
	(0.1367)	(0.0007)
Income Decile 3	-2.7144***	-0.0149***
	(0.1271)	(0.0006)
Income Decile 4	-2.1268***	-0.0120***
	(0.1407)	(0.0007)
Income Decile 5	-1.6887***	-0.0083***
	(0.1509)	(0.0007)
Income Decile 6	-0.9660***	-0.0058***
	(0.1829)	(0.0009)
Income Decile 7	-0.1586	-0.0031***
	(0.2370)	(0.0012)
Income Decile 8	1.0710***	0.0035**
	(0.3077)	(0.0015)
Income Decile 9	3.3427***	0.0164***
	(0.3827)	(0.0019)
Income Decile 10	4.8330***	0.0251***
	(0.4563)	(0.0023)
Observations	2,385,502	2,385,502

NOTE. — All regressions are based on equation (4) and include controls described in Section 5, including worker, CZ, 3-digit-occupation, and 2-digit-occupation-by-year FE. Earnings regressions also include year FE. Robust standard errors in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: Palturai GmbH/Finbot AG (OJV data), 2017—2023; SIAB Version 7523 v1; own calculations.

Appendix References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics* 40, no. S1: 293–340.

Aghion, Philippe, Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard. 2025. How Different Uses of AI Shape Labor Demand: Evidence from France. *AEA Papers and Proceedings* 115: 62–67.

Arntz, Melanie, Michael Böhm, Georg Graetz, Terry Gregory, Florian Lehmer, and Cäcilia Lipowski (2025). *Firm-Level Technology Adoption in Times of Crisis*. IZA Discussion Paper 17846. IZA - Institute of Labor Economics.

Arthur, Rudy. 2021. Studying the UK job market during the COVID-19 crisis with online job ads. *PloS one* 16, no. 5: e0251431.

Ash, Elliott and Stephen Hansen. 2023. Text Algorithms in Economics. *Annual Review of Economics* 15: 659–688.

Bamieh, Omar and Lennar Ziegler (2020). *How Does the COVID-19 Crisis Affect Labor Demand? An Analysis Using Job Board Data From Austria*. IZA Discussion Paper 13801. IZA - Institute of Labor Economics.

Bamieh, Omar and Lennart Ziegler. 2022. Are Remote Work Options the New Standard? Evidence from Vacancy Postings During the COVID-19 Crisis. *Labour Economics* 76: 102179.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. A Practical Guide to Shift-Share Instruments. *Journal of Economic Perspectives* 39, no. 1: 181–204.

Bossler, Mario, Nicole Gürtzgen, Alexander Kubis, Benjamin Küfner, and Martin Popp (2021). *IAB-Stellenerhebung – Version 0018 v1*. Research Data Center (FDZ) of the Institute for Employment Research (IAB).

Büchel, Jan, Vera Demary, Henry Goecke, Enno Kohlisch, Oliver Koppel, Armin Mertens, Christian Rusche, Marc Scheufen, and Jan Wendt (2021). *KI-Monitor 2021: Status quo der Künstlichen Intelligenz in Deutschland*. Gutachten im Auftrag des Bundesverbandes Digitale Wirtschaft (BVDW) e.V. Institut der deutschen Wirtschaft.

Carrillo-Tudela, Carlos, Leo Kaas, and Benjamin Lochner (2023). *Matching through Search Channels*. IZA Discussion Paper 16583. IZA - Institute of Labor Economics.

Chen, Zenan and Jason Chan. 2024. Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise. *Management Science* 70, no. 12: 9101–9117.

Dengler, Katharina and Britta Matthes. 2018. The impacts of digital transformation on the labour market: Substitution potentials of occupations in Germany. *Technological Forecasting and Social Change* 137: 304–316.

Dengler, Katharina, Britta Matthes, and Wiebke Paulus (2014). *Occupational Tasks in the German Labour Market: An alternative measurement on the basis of an expert database*. FDZ Methodology Report 12/2014. Research Data Centre (FDZ) of the Institute for Employment Research (IAB).

Federal Employment Agency (2023). *Employed liable to social security (30 June) [dataset]*. Federal Employment Agency (Germany).

Felten, Edward, Manav Raj, and Robert Seamans. 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42, no. 12: 2195–2217.

Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter. 2006. Imputation Rules to Improve the Education Variable in the IAB Employment Subsample. *Journal of Contextual Economics* 126, no. 3: 405–436.

Forsythe, Eliza, Lisa B. Kahn, Fabian Lange, and David Wiczer. 2020. Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims. *Journal of Public Economics* 189: 104238.

Gathmann, Christina, Felix Grimm, and Erwin Winkler (2024). *AI, Task Changes in Jobs, and Worker Reallocation*. IZA Discussion Paper 17554. IZA - Institute of Labor Economics.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. Text as Data. *Journal of Economic Literature* 57, no. 3: 535–574.

Grienberger, Katharina, Markus Janser, and Florian Lehmer. 2023. The Occupational Panel for Germany. *Journal of Economics and Statistics* 243, no. 6: 711–724.

Gürtzgen, Nicole, Benjamin Lochner, Laura Pohlan, and Gerard J. van den Berg. 2021. Does online search improve the match quality of new hires? *Labour Economics* 70: 101981.

Hensvik, Lena, Thomas Le Barbanchon, and Roland Rathelot. 2021. Job search during the COVID-19 crisis. *Journal of Public Economics* 194: 104349.

Hershbein, Brad and Lisa B. Kahn. 2018. Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review* 108, no. 7: 1737–1772.

Indeed Hiring Lab (2024). *Indeed AI Tracker*. https://data.indeed.com/#/ai. Accessed: 2025-11-16. Indeed.

Landis, J. Richard and Gary G. Koch. 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics* 33, no. 1: 159–174.

Mullahy, John and Edward C. Norton. 2024. Why Transform Y? The Pitfalls of Transformed Regressions with a Mass at Zero. *Oxford Bulletin of Economics and Statistics* 86, no. 2: 417–447.

Rengers, Martina. 2018. Internetgestützte Erfassung offener Stellen: Machbarkeitsstudie im Rahmen eines ESSnet-Projekts zu Big Data. *WISTA - Wirtschaft und Statistik* 5:

Shen, Kailing and Bledi Taska. 2020. Measuring the Impacts of COVID-19 on Job Postings in Australia Using a Reweighting-Estimation-Transformation Approach. *Australian Journal of Labour Economics* 23, no. 2: 153–171.