Beyond Exposure: Predicting AI Adoption Based on Comparative Advantage*

Ilse Lindenlaub
Yale University

Ryungha Oh
University of Chicago

María Alejandra Rodríguez
Yale University

Laura Veldkamp
Columbia University

March 26, 2026

PRELIMINARY AND INCOMPLETE — DO NOT CIRCULATE

Abstract

We study how well widely used occupation-level "AI exposure" measures predict actual AI adoption by workers and propose a new adoption index based on comparative advantage. Using the representative German DiWaBe employee survey linked to worker and establishment information, we compare worker-reported AI use to prominent exposure measures and find that the relationship is weak. Motivated by this gap, we develop a framework in which adoption depends not only on technical feasibility—AI's absolute advantage measured by exposure—but on profitability—AI's comparative (dis)advantage relative to a specific worker—balancing AI productivity against AI user costs and worker productivity against wages. We operationalize this framework at the task level by (i) estimating worker productivity relative to pay, (ii) mapping exposure indices into AI productivity, and (iii) inferring task-specific AI user costs from revealed-preference adoption. The resulting occupation-level index accounts for 60% of cross-occupation variation in observed AI adoption, compared to 14% for an exposure-only model. The two approaches diverge substantially for approximately 30% of workers, highlighting that comparative advantage—not exposure alone—is crucial for assessing AI's labor-market impact.

*The data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and remote data execution.

1 Introduction

Artificial intelligence (AI) has rapidly become one of the most transformative technological inventions of the last decades—hailed by some as a new engine of growth and feared by others as a force that could fundamentally change the workplace. Across academic, policy, and public debates, a common concern has emerged: AI may not merely change how tasks are performed, but in some settings, replace human labor altogether.

To assess the labor market consequences of this technological innovation, a fast-growing body of literature aims to predict the labor market effects of AI relying on occupation-level exposure measures. These exposure measures are summary indices that assign each occupation a predicted "degree of exposure" to AI, typically based on how closely the tasks performed in that job align with what AI systems can plausibly automate, augment, or substitute. In practice, researchers combine information on the task content of occupations (often from standardized databases on job activities) with external indicators of AI capabilities (e.g., which tasks current AI models perform well, or which technologies are being developed and deployed).[1] The resulting exposure scores are then used as proxies for how strongly different jobs are affected by AI, allowing empirical work to compare labor market outcomes—such as employment, wages, hours, and job transitions—across more- versus less-exposed occupations over time. Relying on these occupational exposure measures has become the dominant empirical strategy for predicting the labor market responses to AI: Acemoglu et al. (2022) study employment and wage effects; Georgieff and Hyee (2022) and Pizzinelli et al. (2023) provide cross-country evidence; Hyman et al. (2025) analyze retraining outcomes in exposed occupations; Brynjolfsson et al. (2018) emphasize task reorganization; Freund and Mann (2025) and Althoff and Reichardt (2025) quantify heterogeneous earnings effects using structural models; and Acemoglu (2025) examines aggregate productivity implications.

We find, however, that the empirical link between occupation-level AI exposure and actual AI adoption by workers is weak—the central observation that motivates this paper. We systematically compare widely used AI exposure measures with the actual worker-level AI

[1] Various measures of AI exposure have been proposed, including approaches capturing broad AI-related technologies and occupational characteristics (Webb, 2019; Felten et al., 2021) and more recent measures reflecting advances in generative AI (Eloundou et al., 2024; Handa et al., 2025; Tomlinson et al., 2025).

adoption behavior observed in German data from the DiWaBe survey. DiWaBe is a representative employee survey conducted in 2024 that covers more than 9,000 workers and is linked to establishment-level information. In particular, DiWaBe collects information on respondents' use of AI at work, alongside their occupations, which allows us to directly compare occupation-level AI exposure measures from the literature to individual adoption decisions in our data. To measure occupational exposure to AI, we primarily rely on the task-based measure developed by Eloundou et al. (2024), which quantifies the technical feasibility of performing specific tasks within occupations using AI. This measure aligns well timing-wise with our survey data on actual AI adoption in the DiWaBe since it takes into account the capabilities of large language models. We check robustness with respect to several alternative exposure measures proposed by Felten et al. (2021), Handa et al. (2025), and Webb (2019). Felten et al. (2021) map AI application domains to occupation-specific abilities using expert assessments; Handa et al. (2025) quantify AI adoption using data on interactions with generative AI tools; and Webb (2019) construct exposure based on textual similarity between AI patents and occupational task descriptions.

Figure 1 presents binscatter plots relating actual worker-level adoption rates to the four discussed occupational AI exposure measures. The estimated slopes and thus correlations are modest: Focusing on Panel 1, a 10 percentage point increase in the share of tasks that can be performed by AI according to the exposure measure of Eloundou et al. (2024) is associated with only a 4.6 percentage point increase in the probability of actually adopting AI. We find that this relationship is largely unaffected by the inclusion of controls. More importantly, AI exposure measures explain very little of the variation in adoption behavior: The corresponding regressions yield an R^2 of only 0.008, indicating that occupation-level exposure is a poor predictor of actual AI use by workers. Weak correlations and limited predictive power persist when we consider alternative occupational AI exposure measures (Panels 2–4).

This gap between predicted AI exposure and realized AI adoption is the central motivation for our analysis, since understanding and predicting actual AI adoption is a necessary first step for assessing the labor market effects of this new technology. We examine the sources of this gap and develop a framework for predicting AI adoption that combines expo-

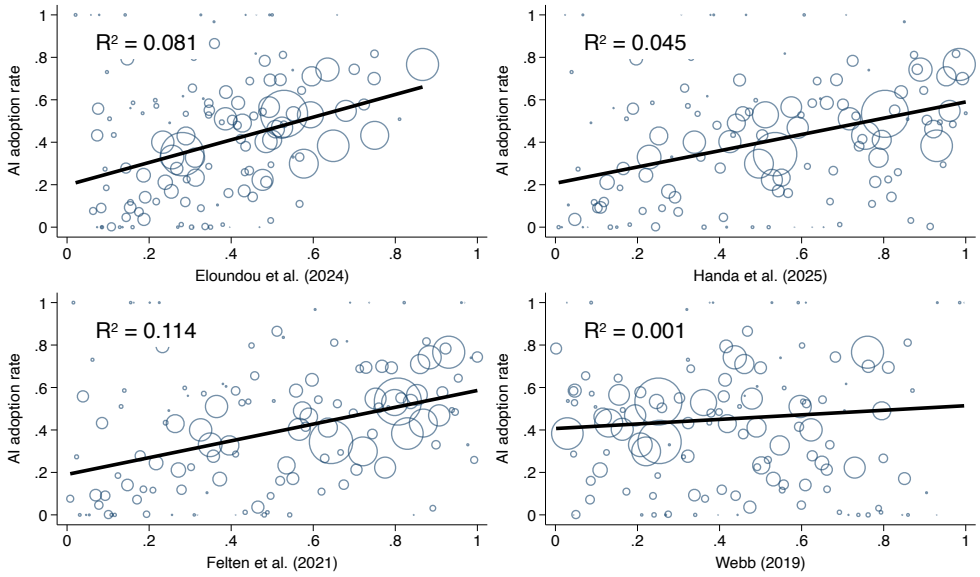


Figure 1: Relation between Actual Adoption by Workers and AI Exposure Measures

Note: The figure relates worker-level AI adoption from DiWaBe to various occupation-level measures of AI exposure. Workers reporting AI use sometimes or more frequently are classified as adopters. We use four alternative exposure measures. For Eloundou et al. (2024) and Handa et al. (2025), exposure is measured as the share of tasks within an occupation exposed to AI; for Felten et al. (2021) and Webb (2019), measures are converted to percentile ranks of occupational AI exposure. Exposure measures defined at the O*NET occupation level are mapped to German KldB 5-digit occupations using occupation crosswalks. Data Sources: DiWaBe for adoption rates and Eloundou et al. (2024), Handa et al. (2025), Felten et al. (2021) and Webb (2019) for exposure measures.

sure measures with worker-level factors and technological costs—margins that jointly shape adoption decisions—and thereby improves predictive accuracy.

The primary contribution of this paper lies in measurement. We start from the premise that AI adoption depends not only on technological feasibility—or the *absolute* advantage in performing a task proxied by AI exposure measures—but also on profitability captured by AI's *comparative* (dis)advantage vis-a-vis a specific worker performing the task. We formalize this idea in a simple framework in which each worker i in occupation o combines a set of task-level outputs, weighted by task importance, to produce overall output. Each task-level output can be produced by either labor or AI according to a linear technology. The worker uses AI for task k if AI has a comparative advantage over labor in producing the task,

$$\frac{Z_{ik}^{\text{worker}}}{w_{io}} < \frac{Z_k^{\text{AI}}}{r_k}, \tag{1}$$

where Z_{ik}^{worker} and Z_k^{AI} denote the productivity of worker and AI, respectively, w_{io} is the wage of worker i, and r_k denotes the user cost of AI services in task k. While tasks that are more exposed to AI—reflected in higher AI productivity Z_k^{AI}—are more likely to be adopted, the adoption decision ultimately hinges on profitability. That is, even when a task is technically feasible for AI (indicated by $Z_k^{\text{AI}} > 0$), adoption may not occur if the user cost of AI is sufficiently high or if production by labor remains more cost-effective. Task-level comparative advantage (1) determines where AI would be used conditional on adoption—the intensive margin. The worker then adopts AI at the extensive margin only if the resulting profit gain, aggregated across those tasks, exceeds the fixed adoption cost.

Our measurement approach proceeds in three steps to recover all components of the adoption condition for each task, enabling us to predict—based on task composition—AI adoption at the occupation level.

In the first step, we start with measuring worker productivity relative to labor cost—the left-hand side of the adoption condition introduced above—using German administrative data. Because task-level worker productivity is not directly observed, we introduce a parsimonious wage-setting framework based on a competitive labor market, in which wages are given by the value marginal product of the worker. Thus, log wages are determined by occupation-specific prices and a weighted average of a worker's task-specific productivities, where task weights capture the relative importance of different tasks within an occupation. Given occupation-specific task weights, we estimate occupation-level wage regressions to recover predicted worker productivity at the task level relative to worker pay, $Z_{ik}^{\text{worker}}/w_{io}$. Intuitively, our approach highlights that this ratio is high—and thus AI adoption becomes less likely—if the worker has a comparative advantage in task k relative to other tasks $k' \neq k$, if pay in the occupation is comparably low, or if the worker's pay is unexpectedly low compared to the worker's systematic productivity.

In a second step, we construct a measure of AI productivity at the task level, Z_k^{AI}, using the widely used occupational exposure indices discussed above. We posit that tasks with greater AI exposure are associated with larger time savings from AI use relative to some reference human, and therefore with higher AI productivity.

In a third step, using structural estimation, we infer the task-specific user cost of AI

r_k (along with moments of the distribution of AI overhead cost) from revealed-preference adoption behavior. Conditional on worker productivity relative to pay as well as AI productivity, variation in AI adoption reflects differences in user costs. We thus recover task-level user costs by finding r_k for each task k to best fit the occupation-level AI adoption shares observed in the data. We are further investigating the extent to which estimated AI user costs depend on task-level determinants such as privacy regulation and safety concerns, as well as on characteristics such as industry and firm size—an analysis that will allow us to characterize the circumstances under which AI adoption is more or less likely.

Our model fits the observed adoption choices well. Based on our preferred goodness of fit measure—a *pseudo* R^2 that captures the share of cross-occupation dispersion in observed AI adoption rates our model is able to account for—our simple model rationalizes almost 60% of the variation of observed adoption behavior. We systematically do model selection, comparing our full model with task-level heterogeneity in AI user cost and worker productivity to nested variants that shut down (i) worker heterogeneity and additionally (ii) user cost heterogeneity implying that only AI productivity/exposure affects adoption. Our model outperforms these nested models by a significant margin: The model that relies only on exposure measures captures 14% of the variation of observed adoption behavior, while the model with added user cost heterogeneity—but without worker heterogeneity—captures 47%, compared to the 60% explained by our full model in which comparative advantage drives adoption.

Our full structural model that predicts AI adoption based on AI profitability thus significantly outperforms the most restricted version that captures only task-level heterogeneity in AI productivity, which we interpret as *AI exposure*. To understand the wedge between exposure and actual AI adoption, we perform a series of diagnostic exercises using the estimated model. We first compare exposure-based measures of AI productivity at the *task level* to AI productivity per unit cost, the object relevant for profitability. The two are not well aligned. For example, "information processing" ranks among the tasks with the highest AI productivity, yet its productivity *per unit cost* falls below that of "coordinating and developing", because its user costs are relatively high. We then examine AI's comparative advantage over labor—the cost-effectiveness of AI relative to workers—which introduces an additional source of divergence. "Information processing" again illustrates this point, as its comparative

advantage is attenuated by relatively high labor productivity in that task. Together, these exercises clarify the economic forces behind the wedge between technical feasibility and realized adoption, and underscore that exposure alone is insufficient for predicting where AI is actually adopted.

We then compare our model's *occupation-level* adoption predictions against those of the exposure-based approach. We classify occupations into four groups based on whether their predicted adoption rate and AI exposure are above or below their respective averages. We find that about 30% of employment, a substantial share, falls into occupations where the two diverge. An analysis relying solely on AI exposure measures would have drawn systematically incorrect conclusions about AI's impact for this large share of workers. This underscores the importance of incorporating AI's comparative advantage over labor for accurately predicting both which workers adopt AI and which workers are most affected by it.

The main outputs of our framework, which we will make available to the research community, include (i) task-level adoption probabilities and their aggregation to occupation-level adoption measures, (ii) an occupation-level adoption-intensity index, and (iii) an occupation-level adoption-gain index. These measures complement existing AI exposure measures by incorporating AI's comparative advantage relative to human labor. They are designed to improve predictions of AI adoption within and across occupations by jointly leveraging information on worker productivity and the cost of technology. Furthermore, we provide a framework for assessing how task-specific advances in AI technologies—through productivity improvements or reductions in user costs—translate into changes in adoption behavior.

In sum, we emphasize that predicting AI adoption from absolute AI productivity in a given task—captured by standard AI exposure measures—provides an incomplete picture. Adoption is inherently a relative decision: It depends on AI productivity relative to the user cost of AI, *compared with* worker productivity relative to labor costs. Both margins—the technology side and the worker side—are central for understanding observed and predicting future adoption patterns.

The Literature. We contribute to the growing literature on the far-reaching impact of AI on workers and jobs.

Most work on AI and labor markets studies how worker outcomes—especially employment and wages—are affected by AI. A relatively small set of papers finds displacement effects. Bonfiglioli et al. (2025) find lower employment growth in more AI-exposed local labor markets, with losses concentrated among lower-skill and production workers. Hui et al. (2024) show that, after the release of ChatGPT, freelancers in more exposed occupations on Upwork experienced declines in employment and earnings. Grennan and Michaely (2020) document negative effects for security analysts, consistent with AI substituting for part of high-skill information processing, while Brynjolfsson et al. (2025a) provide recent payroll-based evidence that any early employment losses are concentrated among young or entry-level workers in highly exposed occupations. By contrast, a larger set of studies finds modest aggregate displacement effects so far. Gathmann et al. (2024) show that in Germany AI changes task content and induces worker reallocation, but employment effects remain small. Humlum and Vestergaard (2025) estimate precise near-zero effects on hours and earnings over the first two years after chatbot adoption in Denmark. Related work emphasizes changes in hiring, task mix, and mobility rather than immediate job destruction: Acemoglu et al. (2022) show that AI adoption changes vacancy composition and reduces non-AI hiring at adopting establishments; Georgieff and Hyee (2022) and Pizzinelli et al. (2023) document substantial cross-country heterogeneity in AI exposure and impacts; and Hyman et al. (2025) study whether workers in exposed occupations can retrain into more AI-intensive jobs.

A second strand studies firms and productivity. Hampole et al. (2025) use employee resume data to proxy firm-level AI adoption and show that AI-exposed tasks have lower labor demand, shifting attention from workers' exposure to firms' technology choices. In parallel, a growing micro literature finds positive but heterogeneous productivity gains from AI across settings. Generative AI raises customer-support productivity in Brynjolfsson et al. (2025c), improves professional writing in Noy and Zhang (2023), boosts performance on many consulting tasks in Dell'Acqua et al. (2026), and speeds software development in Peng et al. (2023) and Cui et al. (2026). Related evidence points to gains in other specialized settings, such as taxi driving Kanazawa et al., 2026, legal work Choi et al., 2024, and radiology Goldsmith-Pinkham et al., 2026.[2] Macro interpretations are more cautious. Acemoglu (2025)

[2]By contrast, Otis et al. (2024) show more uneven effects for Kenyan small-business entrepreneurs based

argues that aggregate gains depend jointly on the share of tasks affected and task-level cost savings, especially in hard-to learn tasks, yielding modest productivity gains over the next 10 years; similarly, Jones and Tonetti (2026) predict that AI will accelerate economic growth but only slowly.

What unifies much of this work is the measurement of the AI shock. *Exposure measures* serve as the empirical proxy for the underlying technological shock in this literature, translating AI capabilities into predicted impacts at the task, occupation, industry, or local labor-market level. Ability-based measures link observed AI progress or applications to O*NET abilities and aggregate these to the occupation level (Felten et al., 2018, 2021; Brynjolfsson et al., 2018; Felten et al., 2023). Patent-based measures infer exposure from overlap between AI patents and job tasks or from industries' use of AI patents (Webb, 2019; Gathmann et al., 2024). LLM-era measures annotate tasks more directly: Eloundou et al. (2024) combine human and GPT-4 assessments to assess task exposure to large language models, while Handa et al. (2025) and Tomlinson et al. (2025) use large-scale usage logs on claude.ai to measure actual usage which they link to occupations. Finally, Bonfiglioli et al. (2025) proxy local exposure with growth in AI-related occupations identified from software requirements in job postings, and Hampole et al. (2025) map AI-related activity in workers' resumes to occupational tasks.

Yet, as we show, exposure—based on technical feasibility of AI—is not the same as adoption, which is based on the profitability of AI. The focus on exposure measures clearly reflects a data constraint: actual adoption is typically unobserved.[3] We overcome this limitation by using worker-level data on actual AI adoption from the German Diwabe survey, linked to administrative data that allow us to measure a wide range of worker and firm characteristics with high precision. Using these data to estimate our structural model of AI adoption, we address an important gap in the literature: the limited evidence on realized AI adoption and its connection to profitability, which is shaped by comparative advantage between workers

on experimental evidence.

[3]Rare exceptions in this literature that also rely on adoption data are Humlum and Vestergaard (2025), who link actual chatbot adoption to Danish administrative data, and Bick et al. (2026), who measure worker-level generative-AI use directly in representative surveys. Both papers, however, have a fundamentally different focus from ours.

and AI.[4]

2 Descriptive Evidence

2.1 The Data

The central input to our analysis is a survey that assesses actual AI adoption of workers in Germany, a dataset we can link to administrative records. We here describe this dataset along with other auxiliary data sources we use in our estimation. Further details on data construction and sample definitions are provided in Appendix D.

AI adoption. Our primary data source is the Digital Transformation and the Changing World of Work (DiWaBe 2.0) survey, a nationally representative employee survey conducted in Germany in 2024. The survey covers 9,835 workers and provides detailed information on technological change at the workplace level, with a particular focus on the diffusion and use of artificial intelligence (AI); see Arntz et al. (2025, 2026) for a detailed description.

A key feature of the DiWaBe survey is that it collects direct information on workers' use of AI at work, and can be linked to rich administrative records, allowing us to combine self-reported AI adoption with high-quality labor market information. We thus have data on individual-level AI adoption as well as aggregated adoption rates to the occupation level. Moreover, the DiWaBe survey builds on an earlier firm-level survey (BIZA II) conducted in 2021, enabling links between worker- and firm-level information. Finally, DiWaBe can be linked to administrative worker records from the Integrated Employment Biographies (IEB), which allows us to obtain information on worker compensation and productivity.

While an earlier wave of DiWaBe was conducted in 2019 (DiWaBe 1.0) and is available for research use (see Müller et al. (2023) for details), the 2024 wave used in this paper is currently confidential but will become available through the Research Data Center (FDZ) at

[4]To our knowledge, the only other paper that attempts to measure the profitability of AI, rather than exposure based on technical feasibility, is Svanberg et al. (2024). Its focus, however, is different. They measure the cost of AI adoption from observables in vision tasks (making up 1.6% of US non-farm compensation) while assuming that technology and humans are equally productive in the exposed tasks. By contrast, we estimate these costs within a structural model of the whole economy allowing for productivity gaps between AI and workers, and systematically compare AI profitability with commonly used exposure measures.

a later point.

AI exposure measures. To measure occupational exposure to AI, we primarily rely on the task-based measure developed by Eloundou et al. (2024), which quantifies the technical feasibility of performing specific tasks within occupations using AI. This measure is particularly well suited for our setting, as it focuses on capabilities of generative AI, closely aligning with the type of AI adoption captured in the DiWaBe 2.0 survey. We assess robustness using several alternative exposure measures proposed by Webb (2019), Felten et al. (2021), and Handa et al. (2025).

Task weights. We adopt the standard view that each occupation consists of various tasks, each of which can be performed either by labor or by AI. To measure the importance of tasks within occupations, we use data from O*NET, the most widely used source of occupational task information and the basis for AI exposure measures.[5] O*NET organizes job content hierarchically from detailed tasks to broader work activities. We focus on clustered Generalized Work Activities (GWAs), resulting in nine task categories in our analysis. To construct task weights, we use the importance scale (IM), which measures how critical each task is for job performance in the occupation. These scales are commonly used in the literature as a measure of how central a descriptor (task, work activity, skill, or ability) is within an occupation (e.g., Blinder, 2007; Firpo et al., 2011; Autor and Dorn, 2013; Vona et al., 2015; Deming, 2017).

Occupations and occupation crosswalk. Because O*NET is defined for U.S. occupations, we need a cross-walk linking it to German occupations. We map German occupation codes (KldB 2010) to O*NET using standard crosswalks via ISCO-08 and SOC classifications. If a German occupation is matched to several U.S. occupations, we aggregate O*NET task weights using simple averages. This way, each of our occupations still has nine tasks and corresponding weights. We use 3-digit occupations for our analysis.

[5]Germany has a comparable survey to O*NET, the Survey of the Working Population on Qualification and Working Conditions in Germany (in short, BIBB), but we prefer the O*NET since it is used as the basis of most exposure measures.

Linked employer-employee data. Additionally, we use the LIAB, a large linked employer–employee panel dataset for Germany provided by the Research Data Centre (FDZ) of the German Federal Employment Agency. The data are based on administrative social security records and contain detailed information on wages and employment histories. We use this dataset to estimate occupation-level wage regressions in order to predict wages and worker productivity in our linked DiWaBe-IEB dataset.[6]

2.2 The Nature of AI Adoption

We begin with summary statistics on the central novelty of our survey: AI adoption choices.

AI adoption by task. Our main variable of interest is AI adoption, constructed from survey answers in the DiWaBe 2.0 on the frequency of AI use across six task types. Table 1 documents the prevalence of AI use across task types. For each task, we report the share of workers who use AI "at least sometimes", which corresponds to our definition of AI adoption. Text generation is the most common use case with an adoption share of 33.4%. Spoken language, diagnostic/analytical, images/videos tasks follow, while cooperative work, and other tasks exhibit much lower adoption rates. Under our baseline definition—a worker is classified as an AI adopter if she uses AI at least sometimes for at least one task—42.3% of workers in the sample have adopted AI. Restricting to workers who use AI "frequently or more" for at least one task yields a more conservative adoption rate of 21.2%.

AI adoption by occupation. While our main analysis exploits variation at the 3-digit occupation level, we first document the broad distribution of AI adoption across 1-digit occupation groups. Table 2 reports AI adoption rates across nine occupation groups. Adoption rates vary substantially across occupations. Natural sciences, geography, and informatics exhibit the highest adoption rate at 60.5%, closely followed by humanities, media, art, culture, and design at 57.4%. Commercial services, sales, and tourism also display a notably high adoption rate of 49.7%. At the other end of the distribution, adoption is lowest in construction, architecture, and technical services and in agriculture, forestry, and farming—at 23.9%

[6] We do not estimate occupation-level wage regressions in the DiWaBe-IEB dataset directly due to its limited size ($\sim$ 9000 observations).

and 18.7%, respectively—a pattern consistent with prior expectations.

Table 1: AI Adoption by Task

Task	AI adoption share (%)
Text	33.42
Spoken language	15.06
Images/videos	11.83
Diagnoses/analyses	12.63
Cooperative work	5.04
Other	6.83

Notes: The AI adoption share is the share of workers using AI at least sometimes for each task type. The frequency scale ranges from never to rarely, sometimes, often, and always. A worker is classified as an AI adopter if she uses AI at least sometimes for at least one task. Each observation is weighted by survey weights. Data Source: DiWaBe 2.0.

AI adoption by industry. Table D.1 (Appendix D.2) reports AI adoption rates across six industry groups. Adoption rates vary substantially across sectors. Knowledge-intensive services exhibit by far the highest adoption rate at 56.9%; within this group, information and communication features the highest adoption rate at 72.3%. Most remaining sectors show intermediate adoption rates between 33% to 43%. By contrast, adoption is lowest in primary industries and utilities, where only 24.9% of workers report using AI.

Table 2: AI Adoption by Occupation Group

Occupation	AI adoption share (%)
Agriculture, forestry, and farming	18.7
Production, raw materials, and manufacturing	43.0
Construction, architecture, and technical services	23.9
Natural sciences, geography, and informatics	60.5
Traffic, logistics, safety, and security	29.2
Commercial services, sales, and tourism	49.7
Business, accounting, law, and administration	44.3
Health care, social sector, teaching, and education	36.6
Humanities, media, art, culture, and design	57.4

Notes: AI adoption share at the 1-digit KldB occupations, weighted by survey weights. Data Source: DiWaBe 2.0.

2.3 Predictive Power of AI Exposure Measures

A growing literature uses occupation-level AI exposure measures as a proxy for the AI shock in order to study the labor market effects of this new technology. We show, however, that these measures have limited predictive power for actual AI adoption, which is our central motivation of proposing a different approach to measure the technology shock.

Occupation-level evidence. We begin by examining the relationship between occupation-level AI exposure and AI adoption. Specifically, we regress occupation-level adoption rates (at the 3-digit KldB level) on standard exposure measures. Table 3 reports the results.

The estimated coefficients are small: a 10 percentage point increase in the share of tasks that can be performed by AI (based on the measure proposed by Eloundou et al., 2024) is associated with only a 4.6 percentage point increase in adoption (column 1). Flexibly allowing for nonlinearities (e.g., higher-order polynomials) increases explanatory power only marginally, indicating that exposure measures explain little of the cross-occupation variation in AI adoption. Alternative AI exposure measures have even more limited explanatory power for actual AI adoption (columns 2-4). Table D.2 replicates these specifications with worker controls (gender, age, education); while the R^2 increases modestly, the limited explanatory power of the exposure measures remains.

A key limitation of existing exposure measures is that they summarize the technical feasibility of AI at the occupation level, but abstract from how profitability—the comparative advantage of AI relative to labor—varies across tasks within the occupation. Motivated by this observation, we examine whether the task composition of an occupation—which tasks are performed and how intensively—contains predictive content for AI adoption beyond what is captured by exposure measures. We find that it does: replacing the scalar exposure measure with task importance weights raises the R^2 from 0.08 to 0.15 at the occupation level. This result suggests that what matters for adoption is not only how exposed an occupation is on average, but also which tasks make up that occupation. Put differently, the composition of tasks within occupations contains information about the distribution of AI profitability across tasks. While reduced-form evidence based on task weights cannot by itself identify the underlying comparative-advantage objects, it is consistent with the view that adoption is

shaped by task-level profitability rather than by technical feasibility at the occupation level alone.

Table 3: AI Adoption and Sources of Variation (Occupation-level)

	AI Exposure Measure				
	(1) Eloundou et al.	(2) Felten et al.	(3) Handa et al.	(4) Webb	(5) Task Weights
AI exposure	0.462***	0.358***	0.297**	-0.031	—
	(0.138)	(0.088)	(0.122)	(0.105)	
Observations	129	129	127	129	129
R^2	0.081	0.114	0.045	0.001	0.152
Corr.	0.285	0.338	0.213	0.026	0.390

Notes: The dependent variable is the occupation-level AI adoption share. Columns (1)–(4) include alternative AI exposure measures constructed by Eloundou et al. (2024), Felten et al. (2021), Handa et al. (2025), and Webb (2019). Column (5) replaces exposure measures with occupation-level task importance weights. Standard errors are reported in parentheses. All variables are constructed at the 3-digit KldB occupation level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data source: DiWaBe 2.0.

Individual-level evidence. We next examine the relationship between AI exposure and adoption at the individual worker level. Table 4 reports regressions of an indicator for AI adoption on exposure measures. The estimated coefficients are small, implying that exposure measures have limited explanatory power for individual adoption, consistent with the occupation-level evidence.

We then augment the specification with worker characteristics. We find that younger, more educated, and male workers are more likely to adopt AI, consistent with Arntz et al. (2026). These variables improve model fit substantially. While the exposure measure of Eloundou et al. (2024) remains a statistically significant predictor, worker heterogeneity accounts for a comparable share of the variation in adoption.

Finally, we incorporate task weights on the right-hand side, as in the occupation-level analysis. This improves predictive performance, with the R^2 increasing from 0.037 to 0.048. Including worker characteristics in this specification further improves fit, indicating that both task structure and worker heterogeneity are important determinants of AI adoption.[7]

Motivated by this evidence, we develop a structural model in which AI adoption depends

[7]The results are qualitatively similar for the alternative exposure measures, although their predictive power is generally lower than that of Eloundou et al. (2024). See Table D.3.

Table 4: AI Adoption and Sources of Variation (Individual-level)

	(1) Eloundou et al.	(2) Eloundou et al.	(3) Task Weights	(4) Task Weights
AI exposure	0.468***	0.445***	—	—
	(0.099)	(0.080)		
Worker controls		Yes		Yes
Observations	6105	6105	6619	6619
R^2	0.037	0.059	0.048	0.064

Notes: The dependent variable is an indicator for worker-level AI adoption. Columns (1) and (2) use the Eloundou et al. AI exposure measure constructed at the 5-digit KldB occupation level. Columns (3) and (4) use occupation-level task importance weights. Standard errors are reported in parentheses. Worker controls include gender, age, squared age, and education (3 groups). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data source: DiWaBe 2.0.

on task-level comparative advantage of AI relative to labor, and occupation-level adoption is obtained by aggregating these task-level profitability terms using occupation-specific task weights. This implies that two occupations with similar average AI productivity—proxied by average AI exposure—may nevertheless exhibit different adoption rates if their task composition differs. We will use this model to decompose the determinants of AI's comparative advantage across tasks into technology and worker components.

3 Conceptual Framework

We develop a simple model, in which AI adoption is guided by its profitability vis-a-vis labor. Our framework is a static partial equilibrium model of AI adoption, in which each occupation achieves a required production level by combining output from various tasks. In turn, each task can be performed by labor or AI: AI is adopted if it lowers the unit costs compared to labor.

3.1 Setting

Workers. The economy is populated by a continuum of workers indexed by i. Each worker is endowed with a vector of task-specific productivities $\{Z_{ik}^{\text{worker}}\}_k$, where $k \in \{1, \ldots, K\}$ indexes tasks, as well as by several observable characteristics such as gender, education etc, which we summarize in vector X_i. Workers are matched to exogenously given production opportunities in occupation o, which is composed of various tasks. Conditional on occupation

o and on the required occupation-level output y, worker i can supply labor across tasks. The labor input l_{ik} denotes labor allocated by worker i to task k, and is a conditional input requirement for producing task output. We assume that worker i earns wage w_{io} for providing her services, which is determined in the background in a competitive labor market and thus a function of the worker task productivities.

Production. Production takes place at the level of production lines within firms. Each production line is associated with a given worker type $\{Z_{ik}^{\text{worker}}, X_i\}$ in occupation o, which needs to render output y_{io}.[8] The price of occupation-level output is denoted by p_o.

We first describe production in the absence of AI adoption. To produce output y_{io}, task-level outputs $\{y_{ik}\}_k$ are combined according to

$$y_{io} = \prod_{k=1}^{K} (y_{ik})^{\alpha_{ok}},$$

where α_{ok} denotes the importance of task k in occupation o, with $\sum_k \alpha_{ok} = 1$. Task-level output is produced using labor according to the linear technology

$$y_{ik} = Z_{ik}^{\text{worker}} l_{ik}.$$

Production may also take place with AI. When adopted, AI augments production but entails an additional multiplicative overhead cost. Specifically, AI adoption introduces a worker-specific cost factor $\exp(\varepsilon_i) \geq 1$, which scales variable costs. The term $\varepsilon_i \geq 0$ captures adoption frictions such as learning, compliance, and workflow integration. We assume $\varepsilon_i \sim F_\varepsilon(\cdot \mid X_i)$, where this distribution can depend on worker characteristics X_i. Formally, let $C_{io}^{\text{AI}}(y)$ denote the minimum variable cost of producing y units when AI is adopted. The effective variable cost is

$$\widetilde{C}_{io}^{\text{AI}}(y) = \exp(\varepsilon_i)\, C_{io}^{\text{AI}}(y).$$

[8] We model AI adoption at the production-unit level, while the data measure adoption at the worker level. Because worker types are defined at a granular level, this distinction is unlikely to be quantitatively important.

With AI adoption, production at the task level is

$$y_{ik} = Z_{ik}^{\text{worker}} l_{ik} + Z_k^{\text{AI}} K_{ik},$$

where Z_k^{AI} denotes task-specific AI productivity and K_{ik} is the AI input used in task k.

3.2 The AI Adoption Choice

For a given required output in a production line of an occupation, y_{io}, the firm chooses how much task output y_{ik} to assign to each task; and for each task, which input supplies that task output at least cost: worker labor or AI. It thereby chooses for each production line whether AI is adopted based on cost considerations.

AI demand at the task level. We start by characterizing the firm's cost minimization problem with and without AI. Detailed derivations are provided in Appendix A. Let D_i be an indicator with $D_i = 1$ if AI is adopted and $D_i = 0$ otherwise.

Absent AI, the firm seeks to allocate labor optimally across tasks to minimize production costs:

$$\min_{l_{ik} \geq 0} w_{io} \sum_k l_{ik} \quad \text{s.t.} \quad \prod_k (Z_{ik}^{\text{worker}} l_{ik})^{\alpha_{ok}} \geq y_{io}. \tag{2}$$

Our model is thus in partial equilibrium: Firms solve for the cost of producing a given output y_{io}, taking prices as given, and we abstract from determining the production scale. We focus on unit costs, as total costs scale linearly with output under constant returns to scale. The unit cost without AI resulting from (2) is

$$UC_{io}^{D_i=0} = \prod_k \left(\frac{w_{io}}{Z_{ik}^{\text{worker}}} \frac{1}{\alpha_{ok}} \right)^{\alpha_{ok}}.$$

With AI, the firm solves

$$\min_{l_{ik} \geq 0, K_{ik} \geq 0} w_{io} \sum_k l_{ik} + \sum_k r_k K_{ik} \quad \text{s.t.} \quad \prod_k (Z_{ik}^{\text{worker}} l_{ik})^{\alpha_{ok}} \geq y_{io} \tag{3}$$

where r_k denotes the unit cost of AI for task k. When AI is available, firms can choose between labor and AI for each task. Define $\mathcal{K}_i^{\text{AI}}$ as the set of tasks where AI is used. When adopting AI, we show that the task-level adoption decision depends on task-level comparative advantage: AI is used in task k if and only if it has lower unit cost than worker i

$$\mathcal{K}_i^{AI} = \left\{ k \in \mathcal{K} : \frac{Z_{ik}^{worker}}{w_{io}} < \frac{Z_k^{AI}}{r_k} \right\}. \tag{4}$$

Due to the linear task output function, the unit cost of producing task-level output is given by the minimum of the two input-specific costs.

In turn, the unit cost of producing output is

$$UC_{io}^{D_i=1} = \prod_k \left(\min \left\{ \frac{w_{io}}{Z_{ik}^{\text{worker}}}, \frac{r_k}{Z_k^{\text{AI}}} \right\} \frac{1}{\alpha_{ok}} \right)^{\alpha_{ok}}.$$

AI adoption. Problem (3) characterizes the 'intensive margin' AI choice. That is, it asks *conditional* on AI being used $D_i = 1$, which tasks is it used for. In turn, on the 'extensive margin', the decision to adopt at all compares the *aggregate* gain from AI to the overhead cost ε_i. That is, AI is adopted if and only if the effective unit cost is smaller, i.e., $UC_i^{D_i=1} \exp(\varepsilon_i) < UC_i^{D_i=0}$. With some rearrangement, this condition leads to the following probability of AI adoption

$$\Pr(D_i = 1) = \Pr\left(\sum_k \alpha_{ok} \max\left\{ \log\left(\frac{Z_k^{\text{AI}}}{r_k} \right) - \log\left(\frac{Z_{ik}^{\text{worker}}}{w_{io}} \right), 0 \right\} > \varepsilon_i \right). \tag{5}$$

Expression (5) has an intuitive interpretation, whereby at the task level AI is adopted based on comparative advantage. Specifically, the term inside the $\max\{\cdot, 0\}$ is a task-level *comparative-advantage gap*. If the gap is positive, AI has lower unit cost than worker i in task k; if the gap is negative, AI is not used in task k and contributes zero to the cost reduction. The occupation weights α_{ok} aggregate these task-level advantages into a log reduction in unit cost of producing occupation output when AI is available.

The corresponding occupation-level adoption rate is

$$s_o := \mathbb{E}[\Pr(D_i = 1)|o], \tag{6}$$

where the expectation is taken over the distribution of worker productivities within occupation o. Equation (6) is the key estimation equation we will fit to the data below.

Nested Models. Our model nests two variants, which we will test against our baseline model that features both heterogeneous worker productivities and user costs across tasks.

The first variant abstracts from the impact of worker heterogeneity, by setting $Z_{ik}^{\text{worker}} = 1$ and assuming a single wage in all occupations, thereby isolating the role of AI-specific factors. Specifically, we consider the following estimation equation:

$$s_o = \Pr\left(\sum_k \alpha_{ok} \max\left\{\log\left(\frac{Z_k^{\text{AI}}}{r_k}\right), 0\right\} > \varepsilon_i\right). \tag{7}$$

On top of this, the second variant shuts down user cost heterogeneity. When also user costs are equal across tasks $r_k = r$, the occupational adoption share simplifies to

$$s_o = \Pr\left(\sum_k \alpha_{ok} \max\left\{\log\left(\frac{Z_k^{\text{AI}}}{r}\right), 0\right\} > \varepsilon_i\right), \tag{8}$$

and is thus purely driven by AI productivity in the various tasks, which we think of as closely related to the discussed exposure measures.

Fitting (7) and (8) to the data and comparing them with (6) will allow us to assess how much predictive power for AI adoption is gained by allowing for heterogeneity in workers and user costs, beyond heterogeneity in AI productivity alone—and thus by accounting for profitability rather than only exposure.

Alternative models. In Appendix B and C, we present several alternative models of AI adoption and show that they yield essentially the same estimation equation as (7). We consider three alternatives. In the first model, we assume the same production structure as in our baseline model but impose a labor constraint that allows us to pin down the scale of

production despite constant returns to scale in production and competitive labor markets, and thus flesh out the full equilibrium. In the second model, we assume that task-level output is produced by Cobb-Douglas function combining labor and non-AI (old) capital. After AI adoption, workers can use AI capital or non-AI capital. Hence, AI is replacing old capital and complementing labor for task-level production. In the third model, we assume labor is the only input; AI adoption increases the labor productivity at the task-level production rather than replacing labor.

All models yield an AI adoption probability of the same form as (7). This shows that our estimation is robust to a variety of different assumptions on the micro-foundations underlying AI adoption, both as far as the functional form and the role of AI in production is concerned.

However, the interpretation of estimates $\theta_k := \log\left(\frac{Z_k^{\text{AI}}}{r_k}\right)$ differ. In our baseline model and the first alternative, it measures the comparative advantage of AI over labor. In the second alternative, it captures the comparative advantage of AI over old capital. In the third alternative, it captures how much AI improves task-level labor productivity.

4 Estimation

Our goal is to estimate the model parameters in order to achieve the best fit between the AI adoption rate implied by the model and the one observed in the data. To do so we proceed in several steps: Outside the model, we will first estimate worker productivity and wages, as well as AI productivity. We will then use the model to estimate AI user costs by fitting occupational adoption shares (6) to the data.

4.1 Worker Productivity and Wages

Estimating Worker Productivity. We first estimate worker productivity across tasks, $\{Z_{ik}^{\text{worker}}\}_k$, which enters the AI adoption decision. We do so outside of the model.

First, we parameterize task-specific productivity as a function of a J-dimensional vector of worker characteristics $X_{it} \equiv \{X_{it,j}\}_{j=1}^{J}$, such that $Z_{ikt}^{\text{worker}} = H_k(X_{it})$. We impose a log-linear

specification:

$$\log H_k(X_{it}) = \sum_j \tilde{\beta}_j^k X_{it,j}.$$

Second, we relate these productivities to observed wages. While in our partial equilibrium theory, there was no need to spell out these details, we assume competitive labor markets with free entry prior to the widespread adoption of AI in the background. As a result, wages w_{iot} equal the value of output $p_{ot}y_{it}$ in the pre-AI period.

We further assume that the mapping from X_{it} to productivity Z_{ik}^{worker} is stable over time, and that relative occupation-level prices are constant, i.e., $p_{ot} = p_o P_t$. Under these assumptions, we estimate the following wage equation,

$$\begin{aligned}\log w_{it} &= \sum_k \alpha_{ok} \log H_k(X_{it}) + \log p_o + \text{FE}_t + \nu_{it} \\ &= \sum_j \beta_j^o X_{it,j} + \text{FE}_o + \text{FE}_t + \nu_{it} \text{ where } \beta_j^o = \sum_k \alpha_{ok}\tilde{\beta}_j^k, \end{aligned} \tag{9}$$

where ν_{it} is a zero-mean measurement error.

Third, we recover the task-level coefficients that determine worker productivity across tasks from the estimated wage regression. That is, we back out $\{\tilde{\beta}_j^k\}$, using the relationship

$$\beta_j^o = \sum_k \alpha_{ok}\tilde{\beta}_j^k,$$

where task weights α_{ok} are treated as known and come from O*NET importance scores as described in Section 2.1. When the task-weight matrix has full column rank, this mapping is invertible. In practice, we estimate $\{\tilde{\beta}_j^k\}$ by choosing the values that best fit this system. This procedure delivers estimates of task-level productivities $\{Z_{ik}^{\text{worker}}\}_k$. See Appendix E.1 for more details.

A worker's comparative advantage in task k relative to other tasks k' is determined by labor productivity for task k per unit labor cost $Z_{ik}^{\text{worker}}/w_{io}$ and, ceteris paribus, affects the likelihood of AI adoption. This ratio is high when worker i is relatively productive in task k compared to other tasks within occupation o, and thus highly productive in k relative to overall compensation. In this case, adopting AI in task k is less likely. Conversely, adoption

is more likely when the worker has relatively low productivity in task k.

Implementation. We use German administrative data from LIAB and pool observations from 2015–2019, prior to the diffusion of generative AI. This choice is motivated by sample size considerations: while the DiWaBe-IEB does contain wage information, its sample contains fewer than 10,000 observations. By contrast, the LIAB provides a substantially larger sample, allowing for a more precise estimation of worker productivities. This is particularly important given the large number of parameters we estimate, including occupation fixed effects in the wage regression. Worker characteristics X_{it} include gender, education (three groups), age and age squared, experience and experience squared, and task exposure. Experience is constructed as the cumulative years of employment observed in the data. Task exposure is constructed as the interaction between years worked in each occupation and task importance, i.e., $\text{exposure}_k = \sum_o (\text{years worked in occupation } o) \cdot \alpha_{ok}$, which captures the effect of task-specific human capital accumulation on pay. We use the estimated wage regression based on LIAB to predict wages in our DiWaBe-IEB sample.[9]

Estimation results. We find substantial heterogeneity in worker productivity. The variance of the log worker productivity, measured by $\sum_k \alpha_{ok} \log Z_{ik}^{\text{worker}}$, equals 0.132. Worker heterogeneity is significant both across and within occupations. The variance of mean log worker productivity across occupations equals 0.05, while the average within-occupation variance equals 0.08, exceeding the across-occupation variation.

Turning to heterogeneity across tasks, worker productivity varies substantially at the task level. The average log productivity gap between the highest-ranked task (identifying objects) and the lowest-ranked task (manual work activities) is approximately 3.7. Furthermore, productivity varies considerably across tasks within workers; the average within-worker variance of log productivity across tasks equals 1.98.

Finally, we find substantial heterogeneity in occupational prices. The variance of $\log p_o$ across occupations is 0.03, indicating an important role for occupational price differences in

[9] To account for wage variation outside the model, we include controls for industry (2-digit), log firm size, firm age and its square, and state fixed effects in the wage regressions, as well as year and month fixed effects. Those variables do not enter the prediction equation though.

wage dispersion.

Construction of model inputs for AI adoption decision. To estimate (6), we construct predicted task productivities and wages for each worker in our DiWaBe-IEB sample,

$$\widehat{\log Z_{ik}}^{\text{worker}} = X_i'\widehat{\hat{\beta}}_k, \qquad \widehat{\log w}_{io} = X_i'\hat{\beta} + \widehat{\log p}_o.$$

We use predicted wages based on the LIAB from 2015-2019 to approximate pre-adoption compensation, which is plausibly determined prior to the widespread use of generative AI adoption.[10]

4.2 AI Productivity

We also construct a measure of AI productivity Z_k^{AI} at the task level outside of our model. Our goal is to capture, for each task k, the amount of task output that AI can produce per unit of AI usage (e.g., per AI-hour).

We map task exposure into productivity using

$$Z_k^{\text{AI}} = \frac{1}{1 - e_k},$$

where $e_k \in [0, 1]$ denotes the exposure of task k from Eloundou et al. (2024). This measure proxies the extent to which AI can reduce the time required to perform task k.[11] Importantly, AI productivity in a task increases in the exposure of the task. We provide a microfoundation for this functional form in Appendix E.2.

By estimating Z_k^{AI} outside of the model, we can in a final step use the model to estimate AI user costs r_k, allowing us to disentangle exposure effects from AI costs in their roles of impacting AI adoption.[12]

[10] Consistent with this assumption, in Figure E.1, we verify empirically that wages remain stable across occupations over time (2015-19 versus 2023) despite heterogeneous changes in AI adoption rates.

[11] Specifically, $e_k = 0$ if there is little to no time saving (less than 50%), $e_k = 0.5$ if AI combined with human input can save at least 50% of the time, and $e_k = 1$ if AI alone can save at least 50% of the time.

[12] Note that our model does not separately identify Z_k^{AI} and r_k, but only their ratio Z_k^{AI}/r_k. Accordingly, our separate construction of Z_k^{AI} serves to disentangle exposure effects from AI costs; but the precise parameterization of Z_k^{AI} is not critical for the estimation fit.

4.3 Estimation Strategy

We estimate the model at both the occupation level–fitting (6)—and the individual level—fitting (5). For both approaches, we assume that overhead cost ε_i follows a truncated normal distribution, and that the cost distribution depends on shifter $C(X_i)$, satisfying $C(X_i) = X_i\gamma$, where X_i includes gender, age, and education (three categories).

Based on observed adoption rates, the model identifies only the relative but not the absolute scale of productivity and cost components. To pin down the level, we normalize the variance of ε_i to 1 and its mean to 10. In the specification without worker heterogeneity, this mean is unconditional. In the extended specification with worker heterogeneity, this normalized mean refers to the reference worker, defined as male, low-educated, and youngest age while the mean overhead costs for other worker groups are shifted by $C(X_i)$. Given this normalization, we estimate a scaling parameter ψ on the labor productivity term, such that $\psi \cdot \log Z_{ik}^{\text{worker}}/w_{ik}$ is identified up to the normalization above, and we subsequently treat this scaled quantity as $\log Z_{ik}^{\text{worker}}/w_{ik}$ throughout.

Occupation-level estimation. We construct observed AI adoption shares at the 3-digit occupation level, denoted by s_o^{data}. Let $s_o(\Theta)$ denote the corresponding model-implied adoption share for occupation o that depends on parameters $\Theta := \{r_1, ..., r_K, \psi, \gamma\}$, obtained by integrating the individual-level adoption probability over the distribution of worker characteristics and shocks; see (6). We estimate the parameter vector Θ by minimum distance. Specifically, we choose Θ to minimize

$$\hat{\Theta} \in \arg\min_{\Theta} Q(\Theta), \quad \text{where} \quad Q(\Theta) = (\mathbf{s}^{\text{data}} - \mathbf{s}(\Theta))'\Omega^{-1}(\mathbf{s}^{\text{data}} - \mathbf{s}(\Theta)).$$

Here, $\mathbf{s}^{\text{data}} = (s_1^{\text{data}}, \ldots, s_O^{\text{data}})'$ and $\mathbf{s}(\Theta) = (s_1(\Theta), \ldots, s_O(\Theta))'$ are vectors of observed and model-predicted adoption shares across occupations. The weighting matrix Ω is diagonal, with elements

$$\Omega = \text{diag}\left(\frac{s_o^{\text{data}}(1 - s_o^{\text{data}})}{W_o}\right),$$

where W_o denotes the sum of sample weights for workers in occupation o.[13]

Individual-level estimation. Let $D_i \in \{0, 1\}$ be an indicator equal to one if worker i adopts AI, and let

$$p_i(\Theta) \equiv \Pr(D_i = 1 \mid Z_{ik}^{\text{worker}}, w_{io}, X_i; \Theta)$$

denote the model-predicted adoption probability; see (5).[14] Assuming conditional independence across workers, the log-likelihood is

$$\ell(\Theta) = \sum_{i=1}^{N} \omega_{io} \left[D_i \log p_i(\theta) + (1 - D_i) \log(1 - p_i(\theta)) \right],$$

where ω_{io} denotes the sampling weight of observation i. We estimate the parameter vector, Θ, by maximizing this log-likelihood:

$$\hat{\Theta} \in \arg\max_{\Theta} \ell(\theta).$$

4.4 The Determinants of AI Adoption

To understand the determinants of AI adoption and assess the relative importance of different mechanisms, we proceed from the most restricted versions of our model to progressively richer specifications, based on the nested models we discussed above. We thus sequentially introduce additional components and evaluate the improvements in model fit.

We begin with an *Exposure* specification that abstracts from heterogeneity in AI costs and worker productivity by imposing $Z_{ik}^{worker} = 1$ and $r_k = r$, so that adoption is driven solely by task-level exposure Z_k^{AI} and task importance weights $\{\alpha_{ok}\}_k$, which corresponds to occupational adoption rate (8). We next introduce task-specific AI costs $\{r_k\}_k$, yielding the *User Cost* specification, in which adoption is determined by the relative advantage of AI across tasks, while workers remain homogeneous, which corresponds to (7).

[13] In practice, we bound s_o^{data} between 0.1 and 0.9 when computing the weighting matrix to prevent any single occupation from receiving disproportionate weight.

[14] At the individual level, we use task shares α_{ok} constructed at the 5-digit (rather than 3-digit) level occupation level since we do not need to worry about sparse cells in the aggregation to occupation-level adoption shares.

We then incorporate heterogeneity in worker productivity $\{Z_{ik}^{\text{worker}}\}_{i,k}$, yielding the *Comparative Advantage* specification. In this specification, workers with a comparative advantage in a given task—i.e., those who are highly productive relative to their pay, compared with AI productivity relative to its cost—are less likely to adopt AI. Finally, we allow for heterogeneous adoption shifters $C(X_i)$ in the overhead cost, capturing differences in the ease of adopting AI across workers, independently of their labor productivity. We refer to this as the *Full* specification, which nests all the previously discussed model variants.

We evaluate model performance using standard measures of fit. For the occupation-level estimation, we assess the predictive performance of each specification using the total variation explained by each model,

$$R^2 = 1 - \frac{\sum_o \Omega_{oo}^{-1} \left(s_o(\hat{\Theta}) - s_o^{\text{data}} \right)^2}{\sum_o \Omega_{oo}^{-1} \left(s_o^{\text{data}} - \bar{s}^{\text{data}} \right)^2},$$

where Ω_{oo}^{-1} denotes the diagonal element of the weighting matrix Ω^{-1}, and $\bar{s}^{\text{data}}$ is the weighted mean of s_o^{data}. We additionally report the correlation $\text{corr}\big(s_o^{\text{data}}, s_o(\hat{\Theta})\big)$. To guide model selection among our nested models, following Andrews and Lu (2001), we compute

$$\text{MAIC} = Q(\hat{\Theta}) + 2N_\Theta, \qquad \text{MBIC} = Q(\hat{\Theta}) + N_\Theta \log O,$$

where $Q(\hat{\Theta})$ is the minimized objective function, O is the number of occupations, and N_Θ is the number of parameters. Lower values indicate a better performance, while balancing increased fit with penalized model complexity.

For the individual-level estimation, we evaluate the model fit using statistics based on the log-likelihood. We compare log-likelihood values across model specifications and conduct likelihood ratio (LR) tests, defined as

$$\text{LR} = -2(\log l_0 - \log l_1),$$

with associated p-values. l_0 and l_1 are the likelihoods of two nested models evaluated at their respective MLE estimates, where l_0 is likelihood under the restricted model and l_1 is the

likelihood under the unrestricted model.

We also compute the McFadden pseudo-R^2, defined as

$$R^2_{\text{McF}} = 1 - \frac{l(\hat{\Theta})}{l_{\text{null}}},$$

where l_{null} is the log-likelihood under the null model, which assigns each individual the average AI adoption rate. Finally, we report standard information criteria,

$$\text{AIC} = -2l(\hat{\Theta}) + 2N_{\Theta}, \qquad \text{BIC} = -2l(\hat{\Theta}) + N_{\Theta} \log W,$$

where $W = \sum_i \omega_{io}$ is the sum of individual weights.

4.5 Estimation Results

We first compare the performance of our nested specifications in predicting actual AI adoption; and then discuss the parameter estimates. Since both the occupation and individual-level estimations yield similar results, we focus on occupation-level estimates in the text and discuss individual-level results in Appendix E.1 as robustness.

Model fit. Table 5 reports the occupation-level fit statistics. The *Exposure* specification delivers weak predictive power for AI adoption at the occupation level, with an R^2 of 0.14 and a correlation between actual and predicted occupational adoption shares of 0.56. This is consistent with the reduced-form evidence in Section 2, where a linear regression of actual adoption rates on the exposure measure yielded similarly modest explanatory power.

Moving to the *User Cost* specification, the model fit improves substantially. The R^2 increases from 0.14 to 0.47 and the correlation between actual and predicted adoption shares rises from 0.56 to 0.69. This highlights that AI adoption is importantly driven by user cost heterogeneity across tasks and thus occupations rather than by AI productivity alone, which is crucially shaped by AI exposure. The sharp improvement in fit indicates that these two objects—exposure/AI productivity and user costs—are not well aligned in the data, a result we return to below.

Allowing for worker heterogeneity in our *Comparative Advantage* specification further

improves the fit from R^2 0.47 to 0.52, and correlation performance measure from 0.69 to 0.72. This suggests that differences in comparative advantage across workers matter for adoption decisions. Workers with a weaker comparative advantage (low Z_{ik}^{worker}/w_{io}) in tasks where AI is more cost effective (high Z_k^{AI}/r_k) are more likely to adopt AI. However, the incremental improvement is modest relative to the gains associated with moving from the *Exposure* to the *User cost* specification. This indicates that task-level variation in AI user costs, interacted with differences in task importance across occupations, is the primary driver of adoption patterns.

Finally, the *Full* specification delivers additional gains in predictive power: the R^2 increases from 0.52 to 0.56, and the correlation measure from 0.72 to 0.75. These improvements are comparable to those from adding labor productivity heterogeneity, suggesting that worker heterogeneity matters both through task-specific comparative advantage and differences in overall adoption propensity.

A different way of assessing performance of the nested models is through model selection criteria, which indeed favor the *Full* Specification, see columns 3 and 4 in Table 5.

Table F.1 (Appendix F.1) reports the results based on the individual-level estimation. The patterns are very similar. The model fit improves as we relax the imposed restrictions, with the largest gains when moving from the *Exposure* to the *User Cost* specification. Moreover, model selection criteria (AIC and BIC) favor the *Full* specification as in the occupation-level estimation. Importantly, our conclusion that task-level variation in AI productivity *relative* to user cost is the central predictor of AI adoption is robust to moving from occupation-level to individual-level estimation. This indicates that the relatively smaller role of worker heterogeneity is not driven by the use of aggregated moments at the occupation level.

Figure 2 visually compares the fit of the *Exposure* specification and the *Full* model by comparing model predicted occupational AI adoption shares to the data counterparts. The full model delivers a substantially better fit, consistent with Table 5. The improvement is most pronounced for occupations with extreme adoption shares—both at the low and high ends of the distribution.

Table 5: Model Fit: Occupation-level Estimation

Model	R^2	Corr.	MAIC	MBIC
Exposure	0.144	0.563	5083	5086
User Cost	0.474	0.689	3139	3164
Comparative Advantage	0.515	0.718	2900	2928
Full	0.561	0.749	2636	2679

Notes: The table reports model fit statistics from the occupation-level estimation. For each specification, we report R^2, the correlation between predicted and actual adoption rates, and model selection criteria (AIC and BIC). Specifications are ordered from the most restricted (Exposure) to the least restricted one (Full).

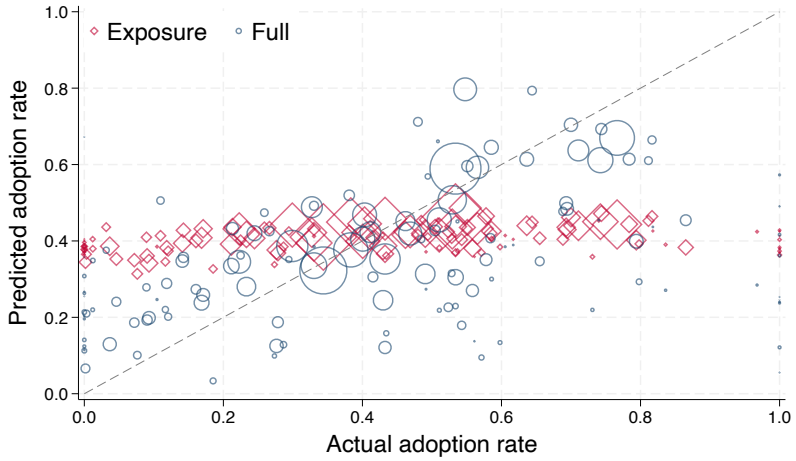


Figure 2: Model Fit for Different Model Specifications: Exposure vs. Full

Notes: The figure plots occupation-level predicted AI adoption rates (vertical axis) against actual adoption rates (horizontal axis) under (i) the *Exposure* specification (red) and (ii) the *Full* model (blue). Each dot represents an occupation and is weighted by the sum of survey weights of workers in that occupation.

Estimated parameters. Table 6 reports the model estimates of task-level AI productivity, costs, and their ratio, as well as parameters related to worker heterogeneity under the full specification.

We find substantial heterogeneity in AI comparative advantage across tasks, indicated by a large variation in $\frac{Z_k^{\text{AI}}}{r_k}$. Cognitive tasks such as reasoning/decision (task 4) and technical activities (task 6) exhibit the highest productivity per unit cost, while manual and communication-intensive tasks display substantially lower values. Quantitatively, the gap between the highest and lowest tasks is large, with top tasks exhibiting nearly twice the productivity per unit cost relative to the bottom ones.

Table 6: Estimates of AI Productivity, User Costs and Worker Heterogeneity

Panel A: Task-level AI productivity and costs

	Tasks	Z_k^{AI}	r_k	$\log\left(\frac{Z_k^{\text{AI}}}{r_k}\right)$
1	Information Input	1.657	2.382	-0.363
				(1.117)
2	Identifying Objects	1.314	0.040	3.504***
				(0.453)
3	Information Processing	2.112	0.029	4.277***
				(0.323)
4	Reasoning/Decision	2.140	0.005	6.150***
				(0.334)
5	Manual Work Activities	1.038	0.095	2.388***
				(0.307)
6	Technical Activities	2.711	0.001	8.388***
				(0.378)
7	Communicating and Interacting	1.544	0.013	4.809***
				(0.299)
8	Coordinating/Developing	1.481	0.013	4.725***
				(0.327)
9	Administering	1.767	0.009	5.282***
				(0.380)

Panel B: Task-level worker heterogeneity

	Parameters	Estimates
ψ	Weight on labor productivity heterogeneity	0.850***
		(0.062)
γ_1	Female	0.513***
		(0.051)
γ_2	Age (/10)	-0.274**
		(0.136)
γ_3	Age squared (/100)	0.047***
		(0.017)
γ_4	Educ: mid	0.130
		(0.084)
γ_5	Educ: high	-0.349***
		(0.098)

Notes: The table reports estimates from the *Full* model specification. The top panel reports task-level AI productivity (Z_k^{AI}), AI user costs (r_k), and their log ratio. The bottom panel reports parameters governing labor productivity heterogeneity (ψ) and heterogeneity in AI adoption unrelated to worker productivity (γ), which shapes $C(X_i)$ and thus the AI overhead cost distribution. Efficient standard errors are reported in parentheses, computed using the Delta method. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Importantly, a stronger comparative advantage does not necessarily arise from higher productivity alone. Tasks with high exposure-based productivity rankings do not always exhibit high productivity per unit cost once costs are taken into account. For example, task 3 (Information Processing) has a higher exposure ranking than task 8 (Coordinating/Developing), indicated by $Z_3^{AI} > Z_8^{AI}$, yet exhibits a weaker comparative advantage $Z_3^{\text{AI}}/r_3 < Z_8^{\text{AI}}/r_8$, because of substantially higher user cost of AI, $r_3 > r_8$.

Turning to worker heterogeneity, our estimates yield intuitive results on how individual characteristics such as age, gender, and education shape AI adoption. In line with existing evidence, as well as our own findings from the DiWaBe survey, these attributes are important predictors of AI adoption, independent of their effect on worker productivity: Younger and highly educated individuals, as well as men, are more likely to adopt AI. Our model captures this through lower means of the AI overhead-cost distributions for these groups.

Finally, the importance of worker heterogeneity across tasks for AI adoption is governed by the estimate of ψ. The standard deviation of the labor productivity per unit cost $\log\left(\frac{Z_{ik}^{\text{worker}}}{w_{io}}\right)$ scaled by ψ equals 1.01, which is of the same order of magnitude as that of the AI productivity per unit cost, suggesting that the two sources of variation are broadly comparable. In Table F.2, we report summary statistics of the labor productivity per unit cost across tasks.

5 AI Adoption Driven by Comparative Advantage

Comparative advantage of AI over labor. Our central finding is that AI adoption is driven not by AI productivity alone, but by AI's comparative advantage over labor, that is, AI productivity per unit cost relative to worker productivity per unit cost. As shown in equation (5), occupations in which AI has a stronger comparative advantage in key tasks are more likely to adopt.

We now report summary statistics of $\log\left(Z_k^{\text{AI}}/r_k\right) - \log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$ for each task in Table 7, capturing the strength of AI's comparative advantage. Overall, the ranking of AI's comparative advantage across tasks closely mirrors the ranking of AI productivity *per unit cost*. For example, the task "technical activities" exhibits both the highest productivity per unit cost and the highest comparative advantage over labor. The task "information

processing", however, departs from this pattern. Despite ranking highly in AI productivity per unit cost, its comparative advantage is attenuated by relatively high labor productivity in that task. By contrast, the task of "communicating and interacting" exhibits relatively low AI productivity but substantially lower user costs, which amplifies AI's comparative advantage in those tasks.

Table 7: Comparative Advantage of AI over Labor by Task

	Task	Mean	Std. Dev.	Q1	Q2	Q3
1	Information Input	-0.633	0.769	-1.125	-0.650	-0.185
2	Identifying Objects	1.609	1.119	0.903	1.636	2.322
3	Information Processing	4.455	0.533	4.190	4.502	4.772
4	Reasoning/Decision	5.786	0.665	5.368	5.885	6.238
5	Manual Work Activities	3.635	0.479	3.341	3.667	3.975
6	Technical Activities	8.647	0.924	8.138	8.597	9.078
7	Communicating and Interacting	4.896	0.554	4.522	4.844	5.237
8	Coordinating/Developing	5.472	0.795	4.961	5.432	5.895
9	Administering	5.159	1.079	4.475	5.123	5.879

Notes: This table reports summary statistics of the comparative advantage of AI over labor, measured by $\log\left(Z_k^{\text{AI}}/r_k\right) - \log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$, for each task.

Exposure versus comparative advantage. In which occupations is AI adoption most prevalent? The answer lies in the interaction between an occupation's task composition and task-level comparative advantage. Since AI's comparative advantage varies across tasks, occupations that are more intensive in tasks in which AI holds a stronger comparative advantage exhibit systematically higher adoption rates, a margin that exposure-based measures fail to capture.

Figure 3 plots predicted AI adoption against the exposure measure. We can classify occupations into four quadrants based on whether their adoption rate and exposure are above or below their respective averages: (i) high exposure and high adoption, (ii) low exposure and low adoption, (iii) high exposure but low adoption, and (iv) low exposure but high adoption. Occupations in groups (i) and (ii) are well predicted by the exposure measure, while those in groups (iii) and (iv) represent systematic deviations driven by variation in AI user costs and worker productivity relative to labor costs across tasks.

The occupation "Accounting, controlling, and auditing" illustrates group (iii) with high

exposure but low adoption: It is intensive in the information-processing task, which exhibits high AI productivity yet a modest comparative advantage of AI over labor and, consequently, lower adoption than exposure alone would predict. In turn, the occupation "Teachers in schools of general education" illustrates group (iv) with low exposure but high adoption: It specializes in the task of Coordinating/Developing where AI productivity is not particularly high; yet AI holds a relatively strong comparative advantage, resulting in greater adoption than exposure-based measures would suggest.

A substantial share of occupations falls into groups (iii) and (iv), together accounting for nearly 30% of employment, which suggests that analyses of AI's economic impact based solely on exposure measures would yield inaccurate predictions for a large fraction of workers. Group (iii) alone accounts for 21.4%, comprising occupations with high exposure—and thus high AI productivity—yet modest comparative advantage of AI over labor and, in turn, relatively lower observed adoption rates. These are precisely the occupations that the existing literature, relying on exposure-based measures, would identify as most affected by AI; yet our framework suggests that the impact of AI on these occupations may be considerably more muted than previously thought.

Gains from AI adoption. Our task-based model, in which AI adoption is driven by the comparative advantage or profitability of this new technology, gives rise to a natural measure of the gains from adoption. We here quantify the gains from adopting AI both at the aggregate and the individual level. Using our model notation, the expected gain for a worker is given by

$$\mathbb{E}[\max\{\log \mathrm{UC}_i^{D_i=0} - \log \mathrm{UC}_i^{D_i=1} - \varepsilon_i,\, 0\}],$$

which is the average log change in the effective unit cost of production.

At the aggregate level, the implied cost savings equal 32.31%, corresponding to an overall productivity increase of approximately 47.72% stemming from AI adoption. The distribution of gains exhibits substantial heterogeneity across workers, with a standard deviation of 0.234 among AI adopters. Moreover, gains are positively correlated with log wages, with a correlation of 0.412, suggesting that AI may have important implications for inequality across workers.

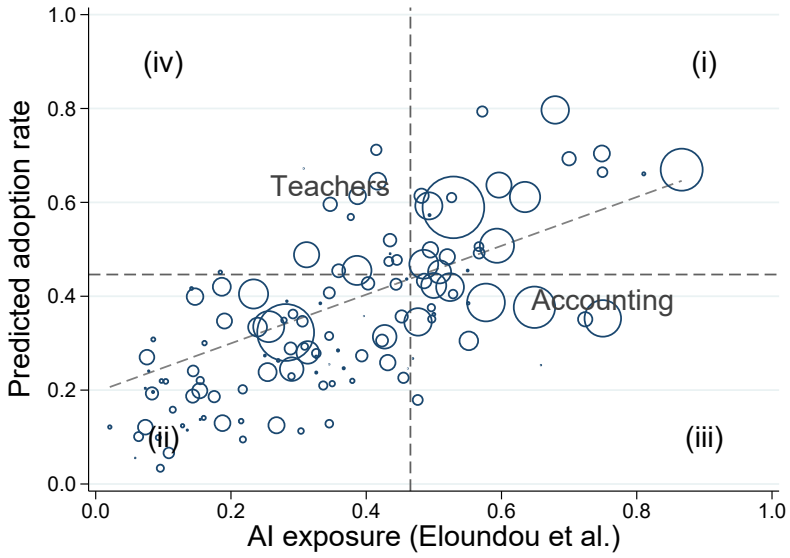


Figure 3: AI Exposure and Predicted Adoption

Notes: The figure plots occupation-level predicted AI adoption under the full specification against the exposure measure developed by Eloundou et al. (2024). Each dot represents an occupation and is weighted by the sum of survey weights of workers in that occupation.

6 Conclusion

In this paper, we show that standard occupation-level AI exposure measures are informative about technical feasibility, but are only weak predictors of realized AI adoption by workers. Using linked German survey worker data linked with administrative records, we can observe actual AI adoption decisions—a rarity in this literature—and document a sizable wedge between exposure and observed adoption. We propose a task-based framework that explains this wedge through comparative advantage: AI is adopted not just based on its productivity—i.e., exposure—but when its productivity per unit cost exceeds worker productivity per unit cost in sufficiently important tasks.

We operationalize this idea in three measurement steps at the task level: (i) we estimate worker task productivity relative to pay from occupation-specific wage regressions; (ii) we map existing exposure indices into task-level AI productivity; and (iii) using the model, we infer task-specific AI user costs from revealed-preference adoption, conditional on worker and AI productivity. Our estimation results highlight that accounting for profitability substantially improves predictive power. A model based only on AI exposure explains little of the

cross-occupation variation in adoption, whereas allowing for heterogeneity in AI user costs and worker productivity yields a much better fit to the data, improving performance by a factor of 4. The largest improvement comes from introducing task-level user-cost heterogeneity, and worker heterogeneity delivers additional explanatory power, underscoring that realized adoption depends on both the technology side and the worker side of the production problem.

More broadly, our findings suggest that predicting the labor-market effects of AI requires moving beyond exposure alone. What matters for adoption is not simply whether AI can do a task, but whether it does so profitably relative to labor. While measuring this comparative advantage is not trivial, it is central for understanding where AI is used today and where it is likely to diffuse next.

References

Acemoglu, D.: 2025, 'The simple macroeconomics of AI'. *Economic Policy* **40**(121), 13–58.

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo: 2022, 'Artificial intelligence and jobs: Evidence from online vacancies'. *Journal of Labor Economics* **40**(S1), S293–S340.

Althoff, L. and H. Reichardt: 2025, 'Task-Specific Technical Change and Comparative Advantage'.

Andrews, D. W. and B. Lu: 2001, 'Consistent model and moment selection procedures for GMM estimation with application to dynamic panel data models'. *Journal of econometrics* **101**(1), 123–164.

Arntz, M., M. Baum, E. Brüll, R. Dorau, M. Hartwig, F. Lehmer, B. Matthes, S. Meyer, O. Schlenker, A. Tisch, et al.: 2025, 'Digital Transformation and the Changing World of Work (DiWaBe 2.0): A Data Source for Research on Artificial Intelligence and Other Technologies in the Workplace'. Technical report, Technical report, 10.21934/BAUA: BERICHT20250319.

Arntz, M., M. Baum, E. Brüll, R. Dorau, M. Hartwig, B. Matthes, S.-C. Meyer, O. Schlenker, A. Tisch, and S. Wischniewski: 2026, 'Low Barriers, High Stakes: Formal and Informal Diffusion of AI in the Workplace'. *ZEW-Centre for European Economic Research Discussion Paper* (26-001).

Arntz, M., K. Dengler, R. Dorau, T. Gregory, M. Hartwig, R. Helmrich, F. Lehmer, B. Matthes, A. Tisch, S. Wischniewski, U. Zierahn, C. Müller, K. Ungerer, and J. Müller: 2023, 'DiWaBe Employee Survey (DIWABE) – Wave 01'.

Autor, D. H. and D. Dorn: 2013, 'The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market'. *American Economic Review* **103**(5), 1553–1597.

Bick, A., A. Blandin, and D. J. Deming: 2026, 'The Rapid Adoption of Generative AI'. *Management Science*. Earlier working-paper version in 2024.

Blinder, A. S.: 2007, 'Offshoring: The Next Industrial Revolution?'. *Foreign Affairs* **85**(2), 113–128.

Bonfiglioli, A., R. Crinò, G. Gancia, and I. Papadakis: 2025, 'Artificial Intelligence and Jobs: Evidence from US Commuting Zones'. *Economic Policy* **40**(121), 145–194.

Brynjolfsson, E., B. Chandar, and R. Chen: 2025a, 'Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence'. *Digital Economy*.

Brynjolfsson, E., D. Li, and L. Raymond: 2025b, 'Generative AI at work'. *The Quarterly Journal of Economics* **140**(2), 889–942.

Brynjolfsson, E., D. Li, and L. R. Raymond: 2025c, 'Generative AI at Work'. *The Quarterly Journal of Economics* **140**(2), 889–942.

Brynjolfsson, E., T. Mitchell, and D. Rock: 2018, 'What can machines learn and what does it mean for occupations and the economy?'. **108**, 43–47.

Choi, J. H., A. B. Monahan, and D. Schwarcz: 2024, 'Lawyering in the Age of Artificial Intelligence'. *Minnesota Law Review* **109**(1), 147–218.

Cui, K. Z., M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz: 2026, 'The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers'. *Management Science*. Published online / forthcoming version circulated in 2025.

Dell'Acqua, F., E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. C. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani: 2026, 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality'. *Organization Science*. Published online March 2026.

Deming, D. J.: 2017, 'The Growing Importance of Social Skills in the Labor Market'. *The Quarterly Journal of Economics* **132**(4), 1593–1640.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock: 2024, 'GPTs are GPTs: Labor market impact potential of LLMs'. *Science* **384**(6702), 1306–1308.

Felten, E., M. Raj, and R. Seamans: 2021, 'Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses'. *Strategic Management Journal* **42**(12), 2195–2217.

Felten, E. W., M. Raj, and R. Seamans: 2018, 'A Method to Link Advances in Artificial Intelligence to Occupational Abilities'. *AEA Papers and Proceedings* **108**, 54–57.

Felten, E. W., M. Raj, and R. Seamans: 2023, 'How Will Language Modelers like ChatGPT Affect Occupations and Industries?'. *arXiv preprint arXiv:2303.01157*.

Firpo, S., N. M. Fortin, and T. Lemieux: 2011, 'Occupational Tasks and Changes in the Wage Structure'. Discussion Paper 5542, Institute of Labor Economics (IZA).

Freund, L. B. and L. F. Mann: 2025, 'Job Transformation, Specialization, and the Labor Market Effects of AI'.

Gathmann, C., F. Grimm, and E. Winkler: 2024, 'AI, Task Changes in Jobs, and Worker Reallocation'. Technical Report 17554, IZA Discussion Paper.

Georgieff, A. and R. Hyee: 2022, 'Artificial intelligence and employment: New cross-country evidence'. *Frontiers in artificial intelligence* **5**, 832736.

Goldsmith-Pinkham, P., C. Tan, and A. K. Zentefis: 2026, 'Human-AI Collaboration in Radiology: The Case of Pulmonary Embolism'. *arXiv preprint arXiv:2601.13379*.

Grennan, J. and R. Michaely: 2020, 'Artificial Intelligence and High-Skilled Work: Evidence from Analysts'. Technical Report 20-84, Swiss Finance Institute.

Hampole, M., D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller: 2025, 'Artificial Intelligence and the Labor Market'. Technical Report 33509, NBER Working Paper.

Handa, K., A. Tamkin, M. McCain, S. Huang, E. Durmus, S. Heck, J. Mueller, J. Hong, S. Ritchie, T. Belonax, et al.: 2025, 'Which economic tasks are performed with ai? evidence from millions of claude conversations'. *arXiv preprint arXiv:2503.04761*.

Hui, X., O. Reshef, and L. Zhou: 2024, 'The Short-Term Effects of Generative AI on Employment: Evidence from an Online Labor Market'. *Organization Science*. Published online September 2024.

Humlum, A. and E. Vestergaard: 2025, 'Large Language Models, Small Labor Market Effects'. Technical Report 33777, NBER Working Paper.

Hyman, B. G., B. Lahey, K. Ni, and L. Pilossoph: 2025, 'How Retrainable Are AI-Exposed Workers?'.

Jones, C. I. and C. Tonetti: 2026, 'Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion'. Unpublished manuscript / slides circulated in 2026.

Kanazawa, K., D. Kawaguchi, H. Shigeoka, and Y. Watanabe: 2026, 'AI, Skill, and Productivity: The Case of Taxi Drivers'. *Management Science* **72**(2), 1376–1388.

Müller, C., K. Ungerer, and J. Müller: 2023, 'DiWaBe employee survey'. FDZ-Datenreport 04/2023 (en), FDZ, Nürnberg.

Noy, S. and W. Zhang: 2023, 'Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence'. *Science* **381**(6654), 187–192.

Otis, N. G., R. P. Clarke, S. Delecourt, D. Holtz, and R. Koning: 2024, 'The Uneven Impact of Generative AI on Entrepreneurial Performance'. Technical report, Harvard Business School Working Paper 24-042.

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer: 2023, 'The Impact of AI on Developer Productivity: Evidence from GitHub Copilot'. *arXiv preprint arXiv:2302.06590*.

Pizzinelli, C., A. J. Panton, M. M. M. Tavares, M. Cazzaniga, and L. Li: 2023, 'Labor market exposure to AI: Cross-country differences and distributional implications'.

Svanberg, M. S., W. Li, M. Fleming, B. C. Goehring, and N. C. Thompson: 2024, 'Beyond AI Exposure: Which Tasks Are Cost-Effective to Automate with Computer Vision?'. Technical report, Working paper.

Tomlinson, K., S. Jaffe, W. Wang, S. Counts, and S. Suri: 2025, 'Working with AI: Measuring the Occupational Implications of Generative AI'. *arXiv preprint arXiv:2507.07935*.

Vona, F., G. Marin, D. Consoli, and D. Popp: 2015, 'Green Skills'. Working Paper 21116, National Bureau of Economic Research.

Webb, M.: 2019, 'The Impact of Artificial Intelligence on the Labor Market'. *Working Paper*.

A Baseline Model: Derivations

A.1 Cost Minimization Problems

The model is solved in two nested stages: (1) task-level: for each task, compute the minimum cost of producing a given y_{ik} with and without AI availability; and (2) occupation-level: given per-task unit costs, minimize total cost of producing occupation output y_{io} under the Cobb-Douglas aggregator.

Stage 1: Task-Level Cost Minimization. Fix worker i, task k, and required task output $y_{ik} > 0$. Without AI adoption, the task problem is

$$\min_{\ell_{ik} \geq 0} \; w_{io}\ell_{ik} \quad \text{s.t.} \quad Z_{ik}^{worker}\ell_{ik} \geq y_{ik}.$$

The constraint binds: $\ell_{ik} = y_{ik}/Z_{ik}^{worker}$. Thus the minimized cost is linear in y_{ik}. Define the labor unit cost

$$c_{ik}^{L} \equiv \frac{w_{io}}{Z_{ik}^{worker}}, \tag{10}$$

which equals the minimized task cost without AI.

When adopting AI, the task problem becomes

$$\min_{\ell_{ik} \geq 0,\, K_{ik} \geq 0} \; w_{io}\ell_{ik} + r_k K_{ik} \quad \text{s.t.} \quad Z_{ik}^{worker}\ell_{ik} + Z_k^{AI} K_{ik} \geq y_{ik}.$$

Generically, the solution is a corner: the technology with the lower unit cost produces the task. Define the AI unit cost

$$c_k^{AI} \equiv \frac{r_k}{Z_k^{AI}}. \tag{11}$$

Then the minimized task cost with AI available is $\min\{c_{ik}^{L}, c_k^{AI}\}\, y_{ik}$. Define the set of tasks in which AI is used when available:

$$K_{io}^{AI} \equiv \left\{k \in \mathcal{K} : c_k^{AI} < c_{ik}^{L}\right\} = \left\{k \in \mathcal{K} : \frac{Z_{ik}^{worker}}{w_{io}} < \frac{Z_k^{AI}}{r_k}\right\}. \tag{12}$$

This is the task-level comparative-advantage condition in (4).

Stage 2: Occupation-Level Cost Minimization Given per-task unit costs, the occupation-level cost-minimization problem for producing $y_{io} = y$ units is, for each regime $d \in \{D_i = 0, D_i = 1\}$,

$$\min_{\{y_{ik} \geq 0\}_{k \in \mathcal{K}(o)}} \sum_{k \in \mathcal{K}(o)} c_{ik}^{d} \, y_{ik} \quad \text{s.t.} \quad \prod_{k \in \mathcal{K}(o)} y_{ik}^{\alpha_{ok}} \geq y,$$

where the relevant unit costs are

$$c_{ik}^{D_i=0} = \frac{w_{io}}{Z_{ik}^{worker}}, \qquad c_{ik}^{D_i=1} = \min\left\{ \frac{w_{io}}{Z_{ik}^{worker}}, \frac{r_k}{Z_k^{AI}} \right\}.$$

Let λ be the multiplier on the production constraint (which binds at optimum). The Lagrangian is

$$\mathcal{L} = \sum_k c_{ik}^{d} y_{ik} + \lambda \left(y - \prod_k y_{ik}^{\alpha_{ok}} \right).$$

FOC for each task k:

$$c_{ik}^{d} = \lambda \alpha_{ok} \frac{\prod_j y_{ij}^{\alpha_{oj}}}{y_{ik}} = \lambda \alpha_{ok} \frac{y}{y_{ik}}, \tag{13}$$

where the second equality uses that the constraint binds, $\prod_j y_{ij}^{\alpha_{oj}} = y$.

Rearranging gives

$$y_{ik} = \lambda \alpha_{ok} \frac{y}{c_{ik}^{d}}.$$

Multiplying (13) by y_{ik} yields

$$c_{ik}^{d} y_{ik} = \lambda \alpha_{ok} y.$$

Summing over k and using $\sum_k \alpha_{ok} = 1$ implies total cost $C_{io}^{d}(y) \equiv \sum_k c_{ik}^{d} y_{ik}$ satisfies

$$C_{io}^{d}(y) = \lambda y,$$

so expenditures satisfy the share rule

$$c_{ik}^{d} y_{ik} = \alpha_{ok} \, C_{io}^{d}(y). \tag{14}$$

Thus, in this Cobb Douglas setting, α_{ok} are both production exponents and cost shares.

Substituting $y_{ik} = \alpha_{ok} C_{io}^d(y)/c_{ik}^d$ from (14) into the binding constraint gives:

$$y = \prod_k \left(\alpha_{ok} \frac{C_{io}^d(y)}{c_{ik}^d} \right)^{\alpha_{ok}} = C_{io}^d(y) \prod_k \left(\frac{\alpha_{ok}}{c_{ik}^d} \right)^{\alpha_{ok}}.$$

Solving,

$$C_{io}^d(y) = y \prod_{k \in \mathcal{K}(o)} \left(\frac{c_{ik}^d}{\alpha_{ok}} \right)^{\alpha_{ok}}.$$

Define unit cost $UC_{io}^d \equiv C_{io}^d(y)/y$. Then

$$UC_{io}^d = \frac{1}{A_o} \prod_{k \in \mathcal{K}(o)} \left(\frac{c_{ik}^d}{\alpha_{ok}} \right)^{\alpha_{ok}}. \tag{15}$$

A.2 AI Adoption Decision

Define the log unit-cost reduction from having AI available (before overhead) as

$$\Delta_{io} \equiv \log UC_{io}^L - \log UC_{io}^{AI}. \tag{16}$$

Taking logs of (15),

$$\log UC_{io}^d = \sum_k \alpha_{ok} \left(\log c_{ik}^d - \log \alpha_{ok} \right).$$

The constant $-\sum_k \alpha_{ok} \log \alpha_{ok}$ cancels between two regimes, so

$$\Delta_{io} = \sum_{k \in \mathcal{K}} \alpha_{ok} \left(\log c_{ik}^{D_i=0} - \log c_{ik}^{D_i=1} \right).$$

Plugging in the unit costs for two regimes,

$$\Delta_{io} = \sum_{k \in \mathcal{K}} \alpha_{ok} \max \left\{ \log \left(\frac{c_{ik}^{D_i=0}}{c_k^{D_i=1}} \right), 0 \right\}.$$

Substitute $c_{ik}^{D_i=0} = w_{io}/Z_{ik}^{worker}$ and $c_k^{D_i=1} = r_k/Z_k^{AI}$:

$$\log \left(\frac{c_{ik}^{D_i=0}}{c_k^{D_i=1}} \right) = \log \left(\frac{w_{io}/Z_{ik}^{worker}}{r_k/Z_k^{AI}} \right) = \log \left(\frac{Z_k^{AI}}{r_k} \right) - \log \left(\frac{Z_{ik}^{worker}}{w_{io}} \right).$$

Thus

$$\Delta_{io} = \sum_{k \in \mathcal{K}(o)} \alpha_{ok} \max \left\{ \log \left(\frac{Z_k^{AI}}{r_k} \right) - \log \left(\frac{Z_{ik}^{worker}}{w_{io}} \right), 0 \right\}. \tag{17}$$

The term inside the max operator is a task-level comparative-advantage gap, which equals the gap between the log of AI output per dollar in task and the log of worker-i output per dollar in task. If the gap is positive, AI has lower unit cost than worker i in task k; If the gap is negative, AI is not used in task k and contributes zero to the cost reduction. The occupation weights α_{ok} aggregate these task-level advantages into a log reduction in unit cost of producing occupation output when AI is available.

With a worker-specific cost factor $\exp(\varepsilon_i)$, the effective unit cost under adoption becomes $\exp(\varepsilon_i) UC_{io}^{D_i=1}$, while without adoption it is $UC_{io}^{D_i=0}$. Adoption is optimal if it reduces cost:

$$UC_{io}^{D_i=0} > \exp(\varepsilon_i)\, UC_{io}^{D_i=1} \quad \Longleftrightarrow \quad \log UC_{io}^{D_i=0} - \log UC_{io}^{D_i=1} > \varepsilon_i.$$

Using (16), we can define the adoption indicator D_i as:

$$D_i = \mathbf{1}\{\Delta_i > \varepsilon_i\}.$$

B Alternative Model with Labor Constraint

B.1 Setting

Workers. Time is static. The economy is populated by a continuum of workers indexed by i. Each worker is endowed with a vector of task-specific productivities $\{Z_{ik}^{\text{worker}}\}_k$ where $k \in \{1, \ldots, K\}$ indexes tasks. Workers choose an occupation $o \in 1, \ldots, O$. For each occupation, worker i draws an idiosyncratic preference shock $\{u_{io}\}_o$ and selects the occupation that maximizes expected utility, $\mathbb{E}[w_{io} + u_o]$, where w_{io} denotes the wage of worker in occupation o, and expectations are taken over the distribution of potential wages and preference shocks.

Production. We first explain the production without AI adoption. A worker i in occupation o combines task-level outputs $\{y_{ik}\}_k$ to produce overall output y_i according to

$$y_i = \prod_{k=1}^{K} (y_{ik})^{\alpha_{ok}},$$

where α_{ok} denotes the importance of task k in occupation o, with $\sum_k \alpha_{ok} = 1$. The price of occupation-level output is denoted by p_o.

The key difference from Section 3 is that each worker is endowed with one unit of time and allocates it across tasks, $\sum_k l_{ik} = 1$.

Task-level output is produced using labor:

$$y_{ik} = Z_{ik}^{\text{worker}} l_{ik}.$$

Workers can adopt AI, which augments production. Prior to making the adoption decision, worker i draws an idiosyncratic AI productivity, $Z_i^{\text{AI}} \sim F_z(\cdot|X_i)$, where X_i denotes a vector of worker and firm characteristics, such as age, gender, industry, and firm size. Conditional on adoption, output becomes

$$y_i = Z_i^{\text{AI}} \cdot \prod_k (y_{ik})^{\alpha_{ok}}.$$

With AI adoption, workers can use both labor and AI inputs in each task. Task-level production is given by

$$y_{ik} = Z_{ik}^{\text{worker}} \, l_{ik} + Z_k^{\text{AI}} \, K_{ik},$$

where Z_k^{AI} denotes task-specific AI productivity and K_{ik} is the amount of AI input allocated to task k.

Labor markets. We take a reduced-form approach to wage determination. Wages are set prior to the AI adoption decision and depend on expected output, $w_i = \mathbb{E}[p_o y_i]\theta_i$. The parameter θ_i captures wedges between marginal product and wages. This formulation nests

both competitive labor market, $\theta_i = 1$, and labor market with frictions, $\theta_i \leq 1$.

B.2 Equilibrium

Labor allocation and AI demand. We begin by characterizing optimal task allocation and AI usage at the worker level. All derivations are provided in Appendix B.3.

Absent AI, worker i in occupation o chooses labor allocation across tasks,

$$\max_{l_{ik}} p_o y_i \quad \text{s.t.} \quad \sum_k l_{ik} = 1,$$

where p_o denotes the output price. Given the Cobb–Douglas structure, workers allocate time proportionally to task weights,

$$l_{ik} = \alpha_{ok}.$$

Substituting back, the value of output equals

$$p_o y_i = p_o \prod_k \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}}. \tag{18}$$

When adopting AI, workers choose both labor allocation and AI inputs

$$\max_{l_{ik}, K_{ik}} p_o y_i - \sum_k r_k K_{ik} \quad \text{s.t.} \quad \sum_k l_{ik} = 1,$$

where r_k denotes the unit cost of AI for task k.

Define $\mathcal{K}_i^{\text{AI}}$ as the set of tasks where AI is used. A key observation is that tasks can be ranked by $\frac{1}{Z_{ik}^{\text{worker}}} \cdot \frac{Z_k^{AI}}{r_k}$, which captures the comparative advantage of AI relative to labor. Hence, optimal AI usage follows a cutoff rule,

$$\mathcal{K}_i^{\text{AI}} = \left\{ k \middle| \frac{1}{(1-\alpha_i^{\text{AI}}) p_o y_i} < \frac{1}{Z_{ik}^{\text{worker}}} \frac{Z_k^{\text{AI}}}{r_k} \right\}, \tag{19}$$

where $\alpha_i^{\text{AI}} = \sum_{k \in \mathcal{K}_i^{\text{AI}}} \alpha_{ok}$ denotes the total task share using AI.

Conditional on $\mathcal{K}_i^{\text{AI}}$, optimal inputs are

$$l_{ik} = \frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}} \mathbf{1}\{k \notin \mathcal{K}_i^{\text{AI}}\}, \quad K_{ik} = \alpha_{ok} \frac{p_o y_i}{r_k} \mathbf{1}\{k \in \mathcal{K}_i^{\text{AI}}\}.$$

As before, workers allocate labor across tasks according to task importance weights, but now only among tasks that use labor. Workers demand AI for tasks that belong to the adoption set. The demand for AI is higher for tasks with larger weights α_{ok} and for tasks with lower unit cost r_k. Substituting into the objective, the output value net of AI user cost is given by

$$\begin{aligned} p_o y_i - \sum_k r_k K_k &= (1-\alpha_i^{\text{AI}}) p_o y_i \\ &= (p_o Z_i^{\text{AI}})^{\frac{1}{1-\alpha_i^{\text{AI}}}} \prod_{k \notin \mathcal{K}_i^{\text{AI}}} \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}} \prod_{k \in \mathcal{K}_i^{\text{AI}}} \left(\alpha_{ok} \frac{Z_k^{\text{AI}}}{r_k}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}}. \end{aligned} \tag{20}$$

AI adoption. Worker i adopts AI if and only if the value under adoption (20) exceeds the value under the no-AI benchmark (18),

$$(p_o Z_i^{\text{AI}})^{\frac{1}{1-\alpha_i^{\text{AI}}}} \prod_{k \notin \mathcal{K}_i^{\text{AI}}} \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}} \prod_{k \in \mathcal{K}_i^{\text{AI}}} \left(\alpha_{ok} \frac{Z_k^{\text{AI}}}{r_k}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}} > p_o \prod_k \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}}. \tag{21}$$

Let $D_i = 1$ if worker i adopts AI and $D_i = 0$ otherwise. After rearranging, the adoption probability can be written as

$$\Pr(D_i = 1) = \Pr\left(G_{io} + \ln Z_i^{\text{AI}} > 0,\ \mathcal{K}_i^{\text{AI}} \neq \emptyset\right) \tag{22}$$

$$\text{where} \quad G_{io} = \sum_{k \in \mathcal{K}_i^{\text{AI}}} \alpha_{ok} \left(\ln\left(\frac{1}{Z_{ik}^{\text{worker}}} \frac{Z_k^{\text{AI}}}{r_k}\right) + \ln p_o + \sum_{k'=1}^{K} \alpha_{ok'} \ln\left(\alpha_{ok'} Z_{ik'}^{\text{worker}}\right) \right)$$

The occupation-level adoption rate is

$$s_o = \mathbb{E}[\Pr(D_i = 1) | o], \tag{23}$$

where the expectation is taken over the distribution of worker productivities within occupation o.

Reformulation of the adoption rule. The adoption condition in (21) involves the set $\mathcal{K}_i^{\text{AI}}$, which in principle requires solving a combinatorial problem over all subsets of tasks. However, this is not necessary; The adoption decision admits an equivalent representation that does not require explicitly characterizing $\mathcal{K}_i^{\text{AI}}$, which greatly simplifies computation. To see this, note that worker i adopts AI if and only if there exists a non-empty set $\mathcal{K}$ such that (21) holds. Equivalently, this condition can be written as

$$G_{io}(\mathcal{K}) + \log Z_i^{\text{AI}} > 0,$$

where $G_{io}(\mathcal{K})$ is defined as in (22), but evaluated for an arbitrary subset $\mathcal{K}$, rather than the optimal adoption set $\mathcal{K}_i^{\text{AI}}$. Importantly, this condition can be rewritten in a separable form across tasks. In particular, the adoption decision is equivalent to

$$\Pr\left(\sum_k \alpha_{ok} \max\left\{\ln \frac{Z_k^{AI}}{r_k} \frac{p_o}{Z_{ik}^{\text{worker}}} + \sum_{k'=1}^{K} \alpha_{ok'} \ln(\alpha_{ok'} Z_{ik'}^{\text{worker}}), 0\right\} + \ln Z_i^{\text{AI}} > 0\right).$$

This representation shows that the adoption decision depends only on tasks for which AI has a productivity advantage, and avoids explicitly solving for $\mathcal{K}_i^{\text{AI}}$. We use this formulation for estimation. It is important to note that this representation is purely auxiliary; the set of tasks for which the expression inside the max operator is positive does not, in general, coincide with the optimal adoption set $\mathcal{K}_i^{\text{AI}}$.

As in the nested model in Section 3, when we set $Z_{ik}^{\text{worker}} = 1$ and assume a single wage and common prices across occupations, the estimation equation essentially simplifies to (7),

$$s_o = \Pr\left(\sum_{k=1}^{K} \alpha_{ok} \max\left\{\ln \frac{Z_k^{AI}}{r_k},\ 0\right\} + \log Z_i^{\text{AI}} > 0\right).$$

Adoption set under zero profit condition. Suppose firms observe Z_i^{AI} prior to hiring and operate under free entry, so that profits are zero. Under this condition, wages coincide with the shadow value of labor, $w_{io} = (1 - \alpha_i^{\text{AI}}) p_o y_i$. This implies that the comparative advantage condition can be expressed directly in terms of wages, rather than output, yielding

a simplified characterization of task-level AI adoption, which is identical to (4),

$$\mathcal{K}_i^{\text{AI}} = \left\{ k \middle| \frac{Z_{ik}^{\text{worker}}}{w_{io}} < \frac{Z_k^{\text{AI}}}{r_k} \right\}.$$

B.3 Derivations

In the case of adopting AI, worker i solves the following Lagrangian problem, choosing $\{K_{ik}, l_{ik}\}$:

$$\mathcal{L} = p_o y_i - \sum_k r_k K_{ik} + \lambda \left(1 - \sum_k l_{ik} \right),$$

where r_k denotes the unit cost of AI and p_o the price of output.

Worker i adopts AI for task k if adoption is profitable. For $k \notin \mathcal{K}_i^{\text{AI}}$, we have $y_{ik} = Z_{ik}^{\text{worker}} l_{ik}$, hence

$$(\partial l_{ik}) \quad p_o \alpha_{ok} \frac{y_i}{l_{ik}} = \lambda \qquad (\partial K_{ik}) \quad p_o \alpha_{ok} \frac{y_i}{l_{ik} Z_{ik}^{\text{worker}}} Z_k^{AI} \leq r_k.$$

For $k \in \mathcal{K}_i^{\text{AI}}$, we have $y_{ik} = Z_k^{\text{AI}} K_{ik}$, hence

$$(\partial l_{ik}) \quad p_o \alpha_{ok} \frac{y_i}{Z_k^{AI} K_{ik}} Z_{ik}^{\text{worker}} \leq \lambda \qquad (\partial K_{ik}) \quad p_o \alpha_{ok} \frac{y_i}{K_{ik}} = r_k.$$

AI adoption set For any $k_1 \in \mathcal{K}_i^{\text{AI}}$ and $k_2 \notin \mathcal{K}_i^{\text{AI}}$, we have

$$\frac{\lambda Z_{k_1}^{AI}}{r_{k_1} Z_{ik_1}^{\text{worker}}} > 1 \quad \text{and} \quad \frac{\lambda Z_{k_2}^{AI}}{r_{k_2} Z_{ik_2}^{\text{worker}}} < 1.$$

Thus, if tasks are ordered such that $\frac{Z_k^{AI}}{r_k Z_{ik}^{\text{worker}}}$ is decreasing in k, then there exists a threshold k^* such that $\mathcal{K}_i^{\text{AI}} = \{1, \ldots, k^*\}$.

From the FOC of labor $\alpha_{ok} p_o \frac{y_i}{y_{ik}} Z_{ik}^{\text{worker}} = \alpha_{ok} \frac{p_o y_i}{l_{ik}} = \lambda$ and constraint $\sum_k l_{ik} = 1$, we get $\lambda = (1 - \alpha_i^{\text{AI}}) p_o y_i$, where the term α_i^{AI} denotes the sum of task weights among tasks using AI. Plugging this value into the above conditions, we have the following AI adoption set:

$$\mathcal{K}_i^{\text{AI}} = \left\{ k \middle| \frac{Z_{ik}^{\text{worker}}}{(1 - \alpha_i^{\text{AI}}) p_o y_i} < \frac{Z_k^{\text{AI}}}{r_k} \right\} \text{ where } \alpha_i^{\text{AI}} = \sum_{k \in \mathcal{K}_i^{\text{AI}}} \alpha_{ok}.$$

To summarize, worker i adopts AI for the tasks in which AI has a comparative advantage over labor, measured by

$$\frac{1}{Z_{ik}^{\text{worker}}}\frac{Z_k^{\text{AI}}}{r_k}.$$

Order tasks k with respect to this ratio so that k_1 denotes the task with the largest value. If $\max_k\left\{\frac{1}{Z_{ik}^{\text{worker}}}\frac{Z_k^{\text{AI}}}{r_k}\right\} > \left(p_o Z_i^{\text{AI}}\prod_k(\alpha_{ok}Z_{ik}^{\text{worker}})^{\alpha_{ok}}\right)^{-1}$, the adoption set is characterized by

$$\underbrace{\frac{1}{Z_{ik_K}^{\text{worker}}}\frac{Z_{k_K}^{\text{AI}}}{r_K} < \cdots <}_{\text{no AI adoption}} \frac{1}{(1-\alpha_i^{\text{AI}})p_o y_i} \underbrace{< \cdots < \frac{1}{Z_{ik_1}^{\text{worker}}}\frac{Z_{k_1}^{\text{AI}}}{r_{k_1}}}_{\text{AI adoption}}.$$

Otherwise, the worker does not adopt AI.

For later, we also compute the output and revenue net of AI capital costs,

$$\begin{aligned} y_i &= \left(Z_i^{\text{AI}}\prod_{k\notin\mathcal{K}_i^{\text{AI}}}\left(Z_{ik}^{\text{worker}}\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}\right)^{\alpha_{ok}}\prod_{k\in\mathcal{K}_i^{\text{AI}}}\left(\frac{\alpha_{ok}p_o Z_k^{\text{AI}}}{r_k}\right)^{\alpha_{ok}}\right)^{\frac{1}{1-\alpha_i^{\text{AI}}}} \\ p_o y_i - \sum_k r_k K_k &= (1-\alpha_i^{\text{AI}})p_o y_i \\ &= (p_o Z_i^{\text{AI}})^{\frac{1}{1-\alpha_i^{\text{AI}}}}\prod_{k\notin\mathcal{K}_i^{\text{AI}}}\left(\alpha_{ok}Z_{ik}^{\text{worker}}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}}\prod_{k\in\mathcal{K}_i^{\text{AI}}}\left(\alpha_{ok}\frac{Z_k^{\text{AI}}}{r_k}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}} \end{aligned}$$

AI adoption. We finally characterize the condition under which a worker adopts AI. Worker i adopts AI if AI adoption increases profit, i.e.,

$$(p_o Z_i^{\text{AI}})^{\frac{1}{1-\alpha_i^{\text{AI}}}}\prod_{k\notin\mathcal{K}_i^{\text{AI}}}\left(\alpha_{ok}Z_{ik}^{\text{worker}}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}}\prod_{k\in\mathcal{K}_i^{\text{AI}}}\left(\alpha_{ok}\frac{Z_k^{\text{AI}}}{r_k}\right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\text{AI}}}} > p_o\prod_k\left(\alpha_{ok}Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}}.$$

C Models with Alternative Production Functions

We also present alternative models that yield the same estimation equation (7) of our restricted baseline model.

Cobb-Douglas production for task-level output. Output is produced by combining task-level outputs, each of which is produced using labor and capital based on a Cobb-Douglas

technology (as opposed to the linear technology in Section 3):

$$y_i = \prod_k (y_{ik})^{\alpha_{ok}}, \quad y_{ik} = (Z_{ik}^{\text{worker}} l_{ik})^{1-\xi} (Z_k^K K_k^K)^{\xi},$$

where Z_k^K is the productivity of (non-AI) capital. Let r_k^K denote the unit price of capital. Workers can either adopt AI or not. If a worker decides to adopt AI, this worker must devote a fraction $c \sim F_c(\cdot)$ of her time to learning how to use AI. Hence, when adopting AI, the worker's total time decreases from 1 to $1-c$ and the labor constraint reads $\sum_k l_{ik} = 1-c$.

This specification implies that labor is allocated across tasks in fixed proportions. Consequently, there is no reallocation of labor across tasks following AI adoption, and hence no notion of AI replacing labor in a subset of tasks.

Solving the profit maximization problem without AI, $\max p_o y_i - \sum_k r_k^K K_k^K - w_{io}$ subject to $\sum_k l_{ik} = 1$, the output is given by

$$y_i(c=0, K_k^K(0); K) = (p_o \xi)^{\frac{\xi}{1-\xi}} \prod_k \left(\alpha_{ok} (Z_{ik}^{\text{worker}})^{1-\xi} \left(\frac{Z_k^K}{r_k^K} \right)^{\xi} \right)^{\frac{\alpha_{ok}}{1-\xi}}.$$

In turn, when worker i adopts AI, labor decreases to $1-c$, and productivity and unit costs change from (Z_k^K, r_k^K) to (Z_k^{AI}, r_k). Workers adopt AI for task k if AI has a comparative advantage over (non-AI) capital, i.e., $\frac{Z_k^K}{r_k^K} < \frac{Z_k^{\text{AI}}}{r_k}$.

The output becomes

$$y_i(c, K_k(c)) = (1-c)(p_o \xi)^{\frac{\xi}{1-\xi}} \prod_k \left(\alpha_{ok} (Z_{ik}^{\text{worker}})^{1-\xi} \left(\max \left\{ \frac{Z_k^{\text{AI}}}{r_k}, \frac{Z_k^K}{r_k^K} \right\} \right)^{\xi} \right)^{\frac{\alpha_{ok}}{1-\xi}}.$$

Then AI adoption increases the revenue net of capital costs, i.e., $p_o y_i(c, K_k(c)) - \sum_k r_k K_k > p_o y_i(c=0, K_k^K(0); K) - \sum_k r_k^K K_k^K$, which is equivalent to output being greater when adopting AI since the revenue net of capital costs equals a constant fraction $1-\xi$ of revenue. Hence, worker i adopts AI if and only if

$$\frac{\xi}{1-\xi} \sum_k \alpha_{ok} \max \left\{ \log \left(\frac{Z_k^{\text{AI}}}{r_k} \right) - \log \left(\frac{Z_k^K}{r_k^K} \right), 0 \right\} + \log(1-c) > 0.$$

This condition motivates equation 7:

$$s_o = \Pr\left(\sum_k \alpha_{ok} \max\{\theta_k, 0\} > \varepsilon\right).$$

The term θ_k represents the comparative advantage of AI over capital for each task, which differs from the interpretation of the CD-linear model.

AI as labor productivity improvement We now consider an alternative role of AI. Task-level output is produced by labor, $y_{ik} = Z_{ik}^{\text{worker}} l_{ik}$, before AI adoption as before. However, we assume that AI adoption increases labor productivity from Z_{ik}^{worker} to $(1+\kappa_k)Z_{ik}^{\text{worker}}$, for example by saving time, instead of fully replacing labor. Workers use AI for task k whenever $\kappa_k > 0$. We continue to assume that AI adoption requires a fraction c of time for learning. Worker i adopts AI if and only if

$$(1-c)p_o \prod_k \left((1+\max\{\kappa_k, 0\})\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}} > p_o \prod_k \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}}.$$

AI adoption occurs if and only if, for some $\epsilon \sim F_\epsilon(\cdot)$,

$$\sum_k \alpha_{ok} \log(1+\max\{\kappa_k, 0\}) > \epsilon.$$

Occupational adoption shares are given by

$$s_o = \Pr\left(\sum_k \alpha_{ok} \log(1+\max\{\kappa_k, 0\}) > \varepsilon\right) = \Pr\left(\sum_k \alpha_{ok} \max\{\theta_k, 0\} > \varepsilon\right)$$

Thus, this condition motivates the same empirical estimation specification as in (7), although the interpretation of estimates θ_k differs; it captures how much AI improves labor productivity.

D Descriptive Evidence: Additional Details

D.1 The Data

This section provides additional details on the construction of the data.

O*NET O*NET is a publicly available database that provides detailed information on U.S. occupations, including task descriptions and task-level ratings. For our analysis, we use the task statements, task ratings, and the work-activity hierarchy to construct occupation-level task weights.

The database organizes job content hierarchically from detailed tasks to broader work activities. We aggregate tasks to the level of Generalized Work Activities (GWAs) and further collapse them by grouping parent categories, yielding a set of nine task clusters used in the analysis.

O*NET reports measures of task intensity, including frequency (FT), importance (IM), and relevance (RT). We use the importance scale (IM), which captures how critical each task is for job performance and is reported as the mean rating on a 1–5 scale for each occupation–task pair. For each occupation o and task k, we use the IM scores as raw weights and normalize them within occupation to construct the task-weight matrix.

Occupation crosswalk. Because O*NET is defined for U.S. occupations, we map German occupation codes (KldB 2010) to O*NET using standard crosswalks through ISCO-08 and SOC classifications. The resulting mapping is many-to-many. We aggregate O*NET task weights across mapped occupations using simple averages.

D.2 Additional Descriptive Evidence

Table D.1: AI Adoption by Industry Group

Industry group	AI adoption share (%)
Primary & Utilities (A, B, D, E)	24.9
Manufacturing (C)	43.1
Construction, Trade & Transport (F, G, H)	37.0
Knowledge-intensive services (J, K, L, M)	56.9
Education & Health (P, Q)	39.2
Other services (I, N, O, R, S, T–U)	32.9

Notes: AI adoption share across six industry groups, constructed from the 1-digit WZ2008 industry classification, weighted by the number of workers in each industry. Data source: DiWaBe 2.0.

Table D.2: AI Adoption and Sources of Variation (Occupation-level)

	AI Exposure Measure				
	(1) Eloundou et al.	(2) Felten et al.	(3) Handa et al.	(4) Webb	(5) Task Weights
AI exposure	0.390**	−0.070	0.236*	0.386***	
	(0.160)	(0.112)	(0.129)	(0.120)	
Observations	129	129	127	129	129
R^2	0.110	0.070	0.098	0.139	0.182

Notes: The dependent variable is the occupation-level AI adoption share. Standard errors are reported in parentheses. AI exposure measures are constructed at the 3-digit KldB occupation level. Worker controls include the average female share, age, age squared, and education group shares (three categories) for each occupation. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data source: DiWaBe 2.0.

Table D.3: AI Adoption and Sources of Variation (Individual-Level)

	AI Exposure Measure				
	(1) Eloundou et al.	(2) Felten et al.	(3) Handa et al.	(4) Webb	(5) Task Weights
Panel A: No controls					
AI exposure	0.468***	0.399***	0.256***	0.157**	
	(0.099)	(0.060)	(0.079)	(0.075)	
Worker controls	No	No	No	No	No
Firm controls	No	No	No	No	No
Observations	6,105	6,105	6,042	6,105	6,619
R^2	0.037	0.053	0.017	0.007	0.048
Panel B: Worker controls					
AI exposure	0.445***	0.403***	0.227***	0.089	
	(0.080)	(0.055)	(0.065)	(0.075)	
Worker controls	Yes	Yes	Yes	Yes	Yes
Firm controls	No	No	No	No	No
Observations	6,105	6,105	6,042	6,105	6,619
R^2	0.059	0.072	0.040	0.028	0.064

Notes: The dependent variable is an indicator for worker-level AI adoption. Columns (1)–(4) use alternative AI exposure measures constructed by Eloundou et al. (2024), Felten et al. (2021), Handa et al. (2025), and Webb (2019), measured at the 5-digit KldB occupation level. Standard errors are reported in parentheses. Worker controls include gender, age, squared age, and education (3 categories). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data source: DiWaBe 2.0.

E Estimation: Additional Details

E.1 Estimation of Worker Productivity

This appendix provides additional details on the estimation of task-specific worker productivity.

From Section 4.1, we have $\beta_j^o = \sum_k \alpha_{ok} \tilde{\beta}_j^k$. Stacking this relationship yields

$$(\tilde{\beta}_j^k)_{J \times K} (\alpha_{ok})_{K \times O} = (\beta_j^o)_{J \times O},$$

where K denotes the number of tasks; O the number of occupations; and J the dimension of X_i.

One approach is to estimate $\{\tilde{\beta}_j^k\}$ directly using interacted regressors:

$$\log w_{it} = \sum_{j,k} \tilde{\beta}_j^k \left(\alpha_{ok} X_{it,j}\right) + \mathrm{FE}_o + \mathrm{FE}_t + \varepsilon_{it}.$$

For $j = 1$, $X_{it,j} = 1$, and thus $\tilde{\beta}_1^k$ and FE_o cannot be separately identified from this regression. Hence, we choose $\tilde{\beta}_1^k$ to best fit the estimated FE_o, and attribute the remainder to $\log p_o$.

Stability of wages. To assess whether AI adoption has affected wages, we compare occupation-level wages between 2015–2019 and 2024. Figure E.1 plots mean log wages at the 3-digit occupation level across the two periods. The observations lie close to the 45-degree line, indicating that, despite heterogeneous changes in AI adoption across occupations, relative wages have remained largely stable. This supports our assumption that wages relevant for hiring decisions have not yet been substantially affected by AI adoption.

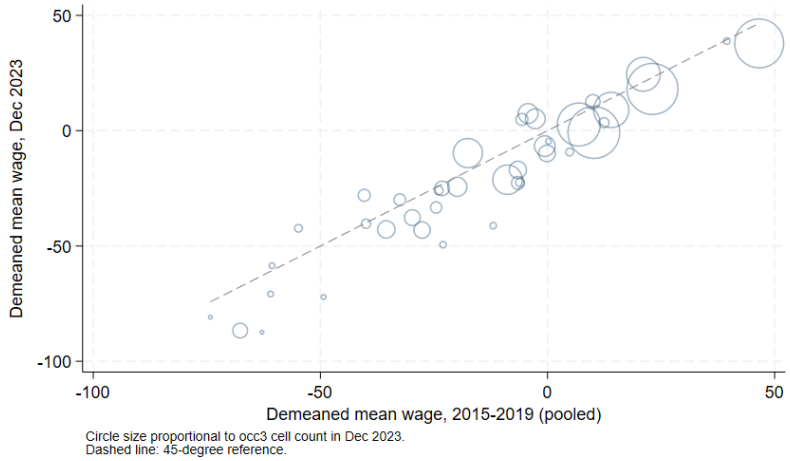


Figure E.1: Wages across Occupations: 2015–2019 vs. 2024

Notes: The figure plots occupation-level mean log wages (demeaned within each period). Marker size is proportional to employment. Some occupations appear clustered due to disclosure restrictions and small sample sizes. Data Sources: DiWaBe 2.0, with wages matched from the IEB.

E.2 AI Productivity

This appendix provides a microfoundation for the mapping from task exposure to AI productivity. We introduce a reference human for each task k with baseline productivity

$$Z_k^H \quad \text{(units of task } k \text{ per hour)}.$$

This serves as a normalization; AI productivity is measured relative to the output of this reference human. Let $\tau_k^H = 1/Z_k^H$ denote the time required for the reference human to complete one unit of task k prior to AI adoption. Suppose AI reduces the time required to perform task k by a fraction Δ_k. Then AI completion time is $\tau_k^{\text{AI}} = (1 - \Delta_k)\tau_k^H$, implying AI productivity

$$Z_k^{\text{AI}} = \frac{1}{\tau_k^{\text{AI}}} = \frac{Z_k^H}{1 - \Delta_k}. \tag{24}$$

We parameterize time savings as $\Delta_k = \kappa e_k$, where $e_k \in [0, 1]$ is the share of tasks that can be performed by AI and $\kappa \in (0, 1]$ scales the effectiveness of AI. Moreover, because the model depends only on relative productivity, i.e., Z_k^{AI} relative to its cost r_k, we can normalize $Z_k^H = 1$ without loss of generality. This gives

$$Z_k^{\text{AI}} = \frac{1}{1 - \kappa e_k}.$$

Our baseline specification sets $\Delta_k = e_k$, where e_k is the AI exposure measure developed by Eloundou et al. (2024).

We consider two alternative approaches. First, if the AI exposure measure is interpreted as the share of tasks that can be replaced by AI, one can calibrate $\kappa < 1$. For example, Acemoglu (2025) uses $\kappa = 0.27$, based on estimates from Noy and Zhang (2023) and Brynjolfsson et al. (2025b). Second, we can use survey-based measures of time savings. Let

$$T_i = 1\{\text{worker reports time savings from AI}\},$$

and define $\bar{T} = \frac{1}{N}\sum_i T_i$. Interpreting $\bar{T}$ as the probability that AI is effective, expected

productivity can be written as

$$Z_k^{\text{AI}} = \bar{T}\frac{1}{1-e_k} + (1-\bar{T}),$$

which, for small e_k, can be approximated by

$$Z_k^{\text{AI}} \approx \frac{1}{1-\bar{T}e_k}.$$

In our survey, the share of workers agreeing that AI has saved time among is $\bar{T} = 34.3\%$.

F Additional Estimation Results

F.1 Individual-Level Estimations

Table F.1: Model Fit: Individual-level Estimation

Model	log-likelihood	LR p-val	McF R^2	AIC	BIC
Exposure	-4449	0.000	0.014	8900	8907
Comparative Advantage	-4348	0.000	0.036	8714	8776
Labor Heterogeneity	-4342	0.000	0.038	8704	8772
Full	-4288	0.000	0.049	8607	8709

Notes: The table reports model fit statistics from individual-level estimation. For each specification, we report the log-likelihood, the likelihood-ratio p-value for the model relative to the null of constant adoption, McFadden's R^2, and model selection criteria (AIC and BIC). Specifications are ordered from the most restricted (Exposure) to the least restricted (Full).

Table F.2: Labor Productivity per Unit Cost

	Task	Mean	Std. Dev.	Q1	Q2	Q3
1	Information Input	0.318	0.905	-0.210	0.338	0.896
2	Identifying Objects	2.229	1.317	1.390	2.197	3.060
3	Information Processing	-0.210	0.627	-0.583	-0.264	0.102
4	Reasoning/Decision	0.429	0.783	-0.103	0.312	0.920
5	Manual Work Activities	-1.467	0.564	-1.868	-1.505	-1.121
6	Technical Activities	-0.305	1.088	-0.812	-0.245	0.294
7	Communicating and Interacting	-0.102	0.652	-0.504	-0.041	0.337
8	Coordinating/Developing	-0.879	0.936	-1.377	-0.832	-0.278
9	Administering	0.144	1.270	-0.703	0.187	0.949

Notes: This table reports summary statistics of labor productivity per unit cost, measured by $\log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$, for each task.

Table F.3: Estimates of AI Comparative Advantage and Worker Heterogeneity

Panel A: Task-level AI productivity and costs

	Tasks	Z_k^{AI}	r_k	$\log\left(\frac{Z_k^{\text{AI}}}{r_k}\right)$
1	Information Input	1.657	0.149	2.408
				(0.561)
2	Identifying Objects	1.314	0.023	4.040
				(0.595)
3	Information Processing	2.112	0.014	5.051
				(0.441)
4	Reasoning/Decision	2.140	0.002	6.836
				(0.443)
5	Manual Work Activities	1.038	0.031	3.498
				(0.357)
6	Technical Activities	2.711	0.004	6.487
				(0.502)
7	Communicating and Interacting	1.544	0.021	4.284
				(0.349)
8	Coordinating/Developing	1.481	0.010	5.005
				(0.388)
9	Administering	1.767	0.008	5.437
				(0.510)

Panel B: Worker heterogeneity

	Parameters	Estimates
ψ	Labor productivity heterogeneity	0.316***
		(0.096)
γ	Female	0.186***
		(0.036)
	Age	-0.080
		(0.135)
	Age squared	0.020
		(0.015)
	Educ: mid	0.058
		(0.054)
	Educ: high	-0.163***
		(0.060)

Notes: The table reports estimates from the full individual-level specification. The top panel reports task-level AI productivity (Z_k^{AI}), costs (r_k), and their log ratio. The bottom panel reports parameters governing labor productivity heterogeneity (ψ) and heterogeneity in AI adoption (γ), which affects $C(X_i)$. Standard errors are reported in parentheses, computed from the numerical Hessian evaluated at the estimates. All estimates in Panel A are statistically significant at the 1% level. Significance stars shown in Panel B: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.