NBER WORKING PAPER SERIES

BEYOND EXPOSURE:
PREDICTING AI ADOPTION BASED ON COMPARATIVE ADVANTAGE

Ilse Lindenlaub
Ryungha Oh
Maria Alejandra Rodriguez
Laura Veldkamp

Working Paper 35271
http://www.nber.org/papers/w35271

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

We thank Christina Patterson for an insightful discussion. We benefited from comments by audiences at the Women in Macro Conference 2026 and the Montreal Workshop on Markets with Frictions 2026. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w35271

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Ilse Lindenlaub, Ryungha Oh, Maria Alejandra Rodriguez, and Laura Veldkamp. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Beyond Exposure: Predicting AI Adoption Based on Comparative Advantage
Ilse Lindenlaub, Ryungha Oh, Maria Alejandra Rodriguez, and Laura Veldkamp
NBER Working Paper No. 35271
May 2026
JEL No. E0, J2, J3

ABSTRACT

We document and explain the gap between measures of AI exposure and measures of AI adoption in the workplace. This leads us to propose a new AI 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 (i.e., AI's absolute advantage measured by exposure) but on profitability (i.e., 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 almost 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.

Ilse Lindenlaub
Yale University
Department of Economics
and NBER
ilse.lindenlaub@yale.edu

Ryungha Oh
The University of Chicago
Booth School of Business
ryungha.oh@chicagobooth.edu

Maria Alejandra Rodriguez
Yale University
Department of Economics
maria.rodriguezvega@yale.edu

Laura Veldkamp
Columbia University
Columbia Business School
and NBER
lv2405@columbia.edu

1 Introduction

Artificial intelligence (AI) is hailed by some as a transformative engine of growth and feared by others as a force that could fundamentally change the workplace. A fast-growing literature predicts the labor market effects of AI, using occupation-level exposure measures.[1] An occupation's predicted "degree of exposure" to AI reflects how closely the tasks performed in that job align with what AI systems can plausibly automate, augment, or substitute. By focusing on what AI is good at, such measures capture the absolute advantage of AI versus human workers. However, trade, labor and growth economists have known for decades that what determines production is not absolute advantage, but comparative advantage. Measuring comparative advantage requires estimating not just the AI ability, but also the worker productivity, worker cost and AI adoption cost. Using an economy-wide representative worker survey, linked to German administrative firm data, we estimate the comparative advantage of AI, across the spectrum of occupations. The resulting comparative advantage index predicts AI adoption substantially better than did the previous exposure measures.

What distinguishes our analysis from others is our ability to estimate comparative advantage of AI. We can do that because of our dataset. It contains information on actual AI adoption by workers—a feature shared by only a small number of existing datasets—and differs from other AI adoption surveys in two important respects: it is representative of the economy as a whole, and survey responses are linked to official establishment-level data. It is the combination of reported AI adoption, paired with a wide array of occupations, and the reliable German administrative records that allows us to infer workers' productivity, their wage and the AI adoption cost. This is what enables us to construct an economy-wide comparative advantage index of workers versus AI and assess its predictive power.

We begin by documenting the moderate correlation of widely used AI exposure measures with the actual worker-level AI adoption behavior. Our adoption data comes from the Digital Transformation and the Changing World of Work (DiWaBe) survey. DiWaBe is a representative employee survey conducted in 2024 that covers more than 9,000 German workers and

[1] Using occupational exposure measures, studies have examined employment and wage effects, retraining outcomes, task reorganization, cross-country patterns, and aggregate productivity implications (Acemoglu et al., 2022; Georgieff and Hyee, 2022; Pizzinelli et al., 2023; Hyman et al., 2025; Brynjolfsson et al., 2018; Freund and Mann, 2025; Althoff and Reichardt, 2025; Acemoglu, 2025).

is linked to establishment-level information and administrative records.

Exposure measures combine information on the task content of occupations from standardized job-activity databases with external indicators of AI capabilities that identify which tasks AI models perform well. 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 takes into account the capabilities of large language models and thereby aligns well with the timing of the DiWaBe adoption survey. We check robustness with respect to several alternative exposure measures: 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 the first panel, 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 5.3 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 relatively little of the observed variation in adoption behavior: The corresponding regressions yield an R^2 of 0.3, 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. To examine the sources of this gap we combine two elements: first, unusual administrative data at the worker level that link actual AI adoption to worker characteristics and wages across all occupations and sectors of the economy; and second, a framework for predicting AI adoption based on both AI exposure as well as technological costs and worker-level factors—margins that jointly shape

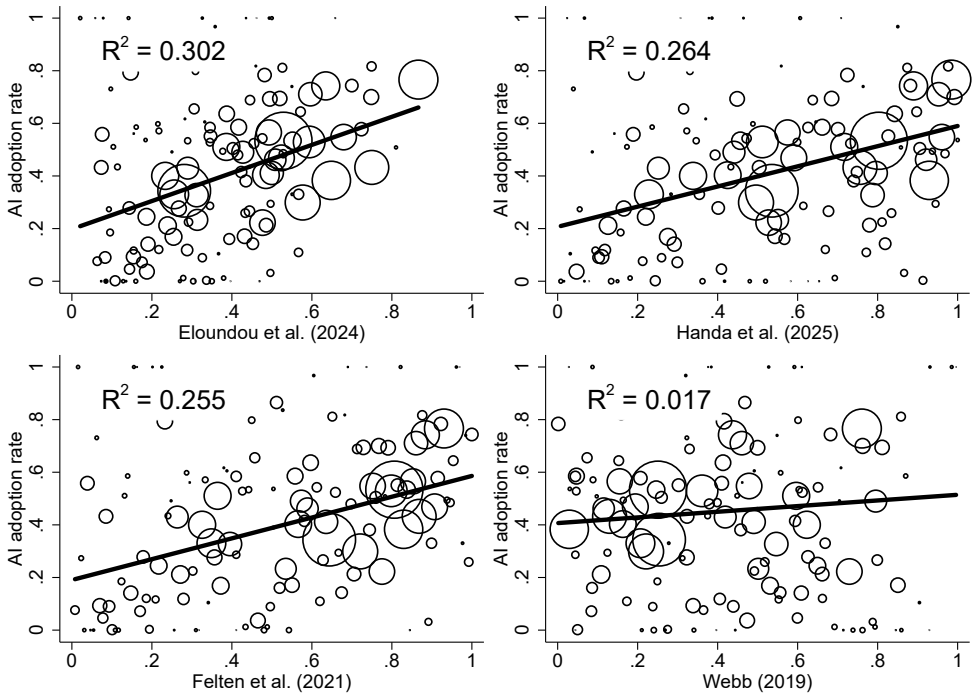


Figure 1: Relation between Actual Adoption by Workers and AI Exposure Measures

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.

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 that serves as our measurement device: 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 overhead cost of adopting AI.

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 $k' \neq k$ or if pay in the occupation is comparably low.

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 identifies 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 develop-

ing", 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 quantify how high human efficiency or high AI user costs in core tasks can attenuate AI's comparative advantage; and show that high AI exposure often fails to predict adoption. Specifically, we classify occupations into four groups based on whether their predicted adoption rate and AI exposure are above or below their respective averages. The occupation "Office clerk" illustrates the first divergent group with high exposure yet low adoption: It is intensive in the information-processing task, which exhibits high AI productivity but also high user costs and human productivity, and therefore a modest comparative advantage of AI over labor. Consequently, adoption is lower than exposure alone would predict. The occupation "Teachers in schools of general education" illustrates the opposite case, an occupation with low exposure but high adoption. This occupation specializes in tasks like coordinating/developing where AI productivity is not particularly high; yet AI holds a relatively strong comparative advantage, stemming from low user costs and worker productivity, which results in greater adoption than exposure-based measures predict. We find that a substantial share of employment—about 30%—is in occupations for which the adoption rates predicted by our model diverge from those implied by AI exposure measures. An analysis relying solely on AI exposure measures would have drawn systematically incorrect conclusions about AI's impact for this large share of workers.

We use the estimated model to project how adoption evolves as AI becomes both more productive and less costly to use. The model predicts rapid aggregate diffusion, with adoption nearly doubling to about 86% over three years when both channels evolve. Decomposing this increase shows that productivity/exposure growth and falling user costs both matter, but that lower user costs play the larger role. At the occupation level, the next wave of adoption

is predicted to come neither from occupations with the lowest current adoption nor from those already near saturation, but from occupations in the middle of the current adoption distribution—examples include accountants and office clerks—where AI is already relevant but substantial room for diffusion remains. For these occupations, adoption thresholds are crossed because AI becomes both more productive and cheaper to use. By contrast, at the lower end of the adoption distribution, adoption gains are driven entirely by falling AI user costs, as AI exposure remains essentially unchanged. In terms of distributional consequences, our model predicts that future AI diffusion will become more spread across the wage distribution, with falling user costs—rather than rising AI exposure—flattening the initially positive adoption–wage gradient.

In sum, we use rare administrative worker-level data that link actual AI adoption to worker characteristics and wages across all occupations and sectors of the economy to understand the mechanism behind adoption and measure the underlying determinants. 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 based on comparative rather than absolute advantage: 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. Our work advances the frontier by directly estimating workers' comparative advantage over AI, allowing us to assess which occupations and sectors are most affected by this technology and which are more insulated from it. We do so by combining three elements: (i) we observe and rationalize actual AI adoption rather than proxying the AI shock via exposure alone; (ii) we recover, from revealed preference, AI user costs across the many tasks of an economy's occupations, rather than directly measuring adoption costs in a narrow setting; and (iii) we measure worker-level productivity and use it—alongside wages and user costs—to compute workers' comparative advantage relative to AI. Previous studies speak to each of these perspectives individually in one way or another, but the combination is what allows us to go beyond exposure and assess adoption profitability.

Our focus is on understanding why AI exposure differs from AI adoption. Occupational exposure measures serve as the empirical proxy for the technology shock in this literature, translating AI capabilities into predicted impacts at the task, occupation, industry, or local labor-market level.[2] Our approach makes use of such pre-existing AI exposure measures as a proxy for absolute AI productivity, specifically LLM-era measures that assign exposure scores directly to tasks: We rely on the measure by Eloundou et al. (2024) who combine human and GPT-4 assessments to assess task exposure to large language models. Similarly, 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. [3]What we add to this is the measurement of worker productivity and AI user costs necessary to understand whether using a worker or adopting AI is more profitable for the task at hand.

The focus on exposure measures clearly reflects a data constraint: actual adoption is typically unobserved. In concentrating on actual adoption by workers, our paper contributes to a nascent literature that, however, has a fundamentally different focus. Humlum and Vestergaard (2025) link chatbot use to Danish administrative data, but focus on a selected subset of occupations and study the effects of AI on worker hours and earnings. Bick et al. (2026a) and Bick et al. (2026b) offer representative coverage of generative AI use across workers, but rely on surveys.[4] Our data combine the strengths of both: We link actual AI adoption to administrative records on worker productivity and wages across all occupations and sectors. This lets us move beyond documenting outcomes toward understanding the mechanism—whether and where AI holds comparative advantage over workers.

[2]Many strands of this literature proxy the AI shock by AI exposure and are adjacent to this work. Jones and Tonetti (2026) focuses on the long-run outcomes, rather than near-term adoption. Related work emphasizes hiring 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; Hyman et al. (2025) study whether workers in exposed occupations can retrain into more AI-intensive jobs. Yet other papers study changes in task mix (Ledingham et al., 2025).

[3]Ability-based exposure measures link observed AI 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). 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.

[4]By contrast, Hampole et al. (2025) focus on *firm* adoption, using employee resume data to proxy firm-level AI adoption and show that AI-exposed tasks have lower labor demand. Other papers exploring firms' predictions related to AI effects are Baslandze et al. (2026) and Bloom et al. (2026).

Measuring whether AI is capable of performing a task better than a human is measuring the absolute advantage of AI.[5] However, just like one country may lack any absolute advantage, but will still produce and trade, workers who lack absolute advantage over AI may still not adopt AI and be displaced. Future employment prospects of workers should depend on their comparative, not their absolute advantage. This requires measurement of the human side—worker productivity relative to wages—as well as the user costs of AI. Recent structural papers also combine task-based comparative advantage with AI exposure measures, but in a different way. They typically proxy AI adoption with exogenous exposure measures rather than deriving it from an endogenous comparative-advantage mechanism. While Acemoglu (2025), Althoff and Reichardt (2025) and Freund and Mann (2025) all have models, in which automation is governed by the comparative advantage of capital versus labor across tasks, in their quantitative exercises, the AI adoption set is not derived from this comparative-advantage condition; instead, this set is taken as exogenous based on AI exposure measures. Our contribution is to derive AI adoption itself from comparative-advantage logic and to measure the primitives that govern this choice.

Predicting AI diffusion based on comparative advantage also requires understanding the friction associated with this new technology, where we emphasize variable, task-level user costs that may be due to compliance and regulation, verification burden or privacy concerns. AI adoption with high productivity gains may still proceed slowly when costs are large, consistent with the gap between rapid advances in AI capabilities and the slower pace of economic diffusion documented in the large-scale expert survey by Karger et al. (2026). Detailed cost estimates of AI technology remain rare, and existing work tends to rely on direct measurement in special settings. Svanberg et al. (2024), for instance, measure the cost of AI adoption by measuring inputs, while focusing on computer vision tasks and assuming equal productivity between technology and humans. We share the motivation to measure AI costs, but take a different approach: rather than measuring costs directly in a specific

[5]Studies of absolute advantage have been used to understand the productivity effects of AI: Generative AI raises customer-support productivity in Brynjolfsson et al. (2025), 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. By contrast, Otis et al. (2024) show more uneven effects for Kenyan small-business entrepreneurs based on experimental evidence.

context, we recover costs from adoption decisions in a structural model of the whole economy.

Many authors study the labor market effects of AI exposure, without uncovering worker comparative advantage.[6] For some questions, just understanding the correlation between exposure and employment is sufficient. However, for making predictions in a changing environment and designing policy we need to understand the mechanisms at work. Specifically, modeling why workers might be replaced allows our AI adoption index to go beyond the concept of exposure. We thereby 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 and AI.

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 briefly along with other auxiliary data sources we use in our estimation. Further details on data construction and sample definitions are provided in Appendix C.

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 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

[6] A handful of studies finds employment impacts of AI in various industries (Bonfiglioli et al., 2025; Hui et al., 2024; Grennan and Michaely, 2020; Brynjolfsson et al., 2025). By contrast, a growing set of studies finds modest aggregate displacement effects so far (Gathmann et al., 2024; Humlum and Vestergaard, 2025).

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. We refer to this linked dataset as DiWaBe-IEB.

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 Centre (FDZ) at 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) (based on patents), Felten et al. (2021) (based on linking common AI applications to workplace abilities and occupations), and Handa et al. (2025) (based on Claude.ai conversations).

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 nearly all existing AI exposure measures.[7] O*NET organizes job content hierarchically from detailed tasks to broader work activities. We focus on clustered Generalized Work Activities (GWA), resulting in nine task categories in our analysis. To construct task weights, we use the importance scale (IM), which is commonly used as a measure of how critical each task is for job performance in the occupation (e.g., Blinder, 2007; Firpo et al., 2011; Autor and Dorn, 2013; Vona et al., 2015; Deming, 2017).

[7]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 using the O*NET since it is used as the basis of most exposure measures.

Occupations and occupation crosswalk. Because O*NET is defined for U.S. occupations, we need a crosswalk 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.

Linked employer-employee data. Additionally, we use the Linked-Employer-Employee-Data of the German Institute of Employment Research (LIAB), which is a large panel dataset for Germany provided by the 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.[8]

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%.

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

Table 1: AI Adoption by Task

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

Notes: The AI adoption share is the fraction 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 occupation. While our main analysis exploits variation at the 3-digit occupation level, we first document the broad distribution of AI adoption across nine 1-digit occupation groups; see Table 2. 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% and 18.7%, respectively—a pattern consistent with prior expectations.

AI adoption by industry. Appendix Table C.1 reports AI adoption rates across six industry groups. Adoption rates vary substantially also 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%. At the other extreme, adoption is lowest in primary industries and utilities, where only 24.9% of workers report using AI.

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—our central motivation

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: The AI adoption share is the fraction of workers using AI at least sometimes for at least one task of the occupation. We report AI adoption share at the 1-digit KldB occupation level, weighted by survey weights. Data Source: DiWaBe 2.0.

for 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 positive but of moderate size: 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 5.3 percentage point increase in adoption (column 1). Moreover, an R^2 of 0.302 indicates that AI exposure has limited explanatory power for AI adoption. Flexibly allowing for nonlinearities (e.g., higher-order polynomials) increases explanatory power only marginally, substantiating 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).

Individual-level evidence. We next examine the relationship between AI exposure and adoption at the individual worker level. Appendix Table C.2 reports regressions of an indicator for AI adoption on exposure measures. The estimated coefficients are broadly in line with those obtained at the occupation level, implying that moving along the distribution of AI exposure translates into only modest changes in individual adoption probabilities. Thus, exposure captures little of the economically relevant variation in adoption behavior, consistent

Table 3: AI Adoption and AI Exposure (Occupation-level)

	(1) Eloundou et al.	(2) Felten et al.	(3) Handa et al.	(4) Webb
AI exposure	0.534***	0.387***	0.359***	0.108
	(0.111)	(0.086)	(0.060)	(0.135)
Observations	129	129	127	129
R^2	0.302	0.255	0.264	0.017
Corr.	0.550	0.505	0.513	0.130

Notes: The dependent variable is the occupation-level AI adoption share. Columns (1)–(4) use alternative AI exposure measures from Eloundou et al. (2024), Felten et al. (2021), Handa et al. (2025), and Webb (2019), respectively. Regressions are weighted by the survey-weighted occupation size N_o^w. Robust standard errors 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.

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.[9]

This evidence highlights a key limitation of existing exposure measures: they capture the technical feasibility of AI at the occupation level, but not its task-level profitability relative to labor. Our reduced-form results suggest that adoption is shaped by this profitability margin, not by technical feasibility alone. We therefore develop a structural model in which AI adoption depends on AI productivity—which standard exposure measures seek to proxy—relative to user costs, compared to worker productivity relative to wages. We use it to decompose 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

[9]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 Appendix Table C.2.

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}.[10] The price of occupation-level output is denoted by p_o.

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 $\sum_k \alpha_{ok} = 1$.

First consider production in the absence of AI adoption. Task-level output is produced

[10] We model AI adoption at the level of a production unit associated with a worker type $(Z_i^{\text{worker}} k, X_i)$ in occupation o, whereas the data record AI use at the worker level. Because worker types are defined at a granular level and because the adoption condition depends on unit-cost comparisons rather than production scale, the production-unit adoption decision maps closely into the observed worker-level adoption choice.

using labor according to the linear technology

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

Production may also take place with AI. In this case, 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.

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

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

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: labor or AI. It thereby chooses for each production line whether AI is adopted based on cost considerations. Detailed derivations are provided in Appendix A.

AI demand at the task level. We start by characterizing the firm's cost minimization problem with and without AI. 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 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 variable 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

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

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, implying

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

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—which is again our focus since total variable costs scale linearly with 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 at the production line and occupation level. Problem (3) characterizes the 'intensive margin' AI choice. That is, it asks *conditional* on AI being used $D_i = 1$, which tasks it is 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 in a production line

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

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

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

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. As a convenient short form, we will sometimes use $\theta_k := \log\left(Z_k^{\mathrm{AI}}/r_k\right)$.

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

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

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

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

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

and is 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 of this new technology rather than only exposure.

Alternative micro-foundations render the same AI adoption share. In Appendix B, we present several alternative micro-foundations underlying the AI adoption choice and show that they yield essentially the same estimation equation as our baseline model. We consider three alternatives. In the first model, we maintain the same production structure as in the baseline but model AI adoption as a worker-level decision subject to a labor constraint, which pins down the scale of production. In the second model, we assume that AI adoption augments labor productivity at the task level rather than automating labor. In the third 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 replaces old capital but complements labor in the task-level production, rather than substituting it. The second and third models are particularly important in light of the ongoing debate over whether AI automates labor, as in our baseline model, or augments it.

All models imply an AI adoption probability of the same form as (5), or a nested version thereof, which we can and will assess through model selection below. This shows that our estimation is robust to a range of assumptions about the micro-foundations of AI adoption, including the functional form of production and the role of AI in production, and to whether output is pinned down in equilibrium or held fixed.

However, the interpretation of estimates θ_k differs. In our baseline model and the first alternative, $\theta_k = \log\left(Z_k^{\mathrm{AI}}/r_k\right)$ and so it measures the comparative advantage of AI over labor. In the second alternative, θ_k captures how much AI augments task-level labor productivity. And in the third alternative, it captures the comparative advantage of AI over old capital.

4 Estimation

Our goal is to estimate the model parameters that achieve the best fit between the AI adoption rate implied by the model and the one observed in the data. 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. The first determinant of AI adoption we estimate is worker productivity across tasks, Z_{ik}^{worker}. We do so outside of the model.

First, we parameterize task-specific productivity as a function of a $\tilde{J}$-dimensional vector of worker characteristics $\tilde{X}_{it}$, such that $Z_{ikt}^{\text{worker}} = H_k(\tilde{X}_{it})$, where we impose a log-linear specification:

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

Second, we relate these productivities to observed wages. The baseline adoption model takes wages as given and does not require us to detail the wage determination. In estimation, however, task-specific worker productivity is unobserved. We therefore use pre-AI wage variation to discipline the mapping from worker characteristics to task productivity. The microfoundation is a competitive pre-AI benchmark with free entry. Under free entry (and CRS in production), price equals unit cost, so the wage per unit of labor input satisfies $w_{it} = p_{ot} \prod_k \left(\alpha_{ok} Z_{ikt}^{\text{worker}}\right)^{\alpha_{ok}}$.[11]

Assuming that the mapping from $\tilde{X}_{it}$ to task productivity is stable over time and that

[11] The wage equation can be derived from pre-AI unit-cost pricing. Let $L_i = \sum_k \ell_{ik}$ and $s_{ik} = \ell_{ik}/L_i$, with $\sum_k s_{ik} = 1$. Pre-AI output is $y_{io} = \prod_k \left(Z_{ik}^{\text{worker}} \ell_{ik}\right)^{\alpha_{ok}} = L_i \prod_k \left(Z_{ik}^{\text{worker}} s_{ik}\right)^{\alpha_{ok}}$. For any given L_i, the optimal task allocation solves $\max_{\{s_{ik}\}} \sum_k \alpha_{ok} \log s_{ik}$ subject to $\sum_k s_{ik} = 1$, implying $s_{ik} = \alpha_{ok}$. Hence $y_{io} = A_{io} L_i$, where $A_{io} = \prod_k (\alpha_{ok} Z_{ik}^{\text{worker}})^{\alpha_{ok}}$. Unit cost is therefore w_{io}/A_{io}. Free entry implies $p_o = w_{io}/A_{io}$, or $w_{io} = p_o A_{io} = p_o \prod_k (\alpha_{ok} Z_{ik}^{\text{worker}})^{\alpha_{ok}}$. Taking logs gives $\log w_{io} = \log p_o + \sum_k \alpha_{ok} \log Z_{ik}^{\text{worker}} + \sum_k \alpha_{ok} \log \alpha_{ok}$, where the final term is absorbed by the occupation fixed effect.

relative occupation prices are constant, $p_{ot} = p_o P_t$, this yields the estimating equation

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

and 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

$$\beta_j^o = \sum_k \alpha_{ok} \tilde{\beta}_j^k, \tag{10}$$

where task weights α_{ok} come from O*NET importance scores and are treated as known (see Section 2.1).

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 estimated to be 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 a relatively low productivity in task k.

Identification. When the task-weight matrix, $A := [\alpha_{ok}]$, has full column rank, mapping (10) is invertible and we can thus identify $\{\tilde{\beta}_j^k\}$ and thus task-level productivities $\{Z_{ik}^{\text{worker}}\}_k$ based on wage regression 9. In practice, we estimate $\{\tilde{\beta}_j^k\}$ by choosing the values that best fit system (10). This procedure delivers estimates of $\{Z_{ik}^{\text{worker}}\}_k$. See Appendix E.1 for details.

Implementation. We use German administrative data from LIAB to estimate the wage regression. 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 using an occupation-level wage regression. This is particularly

important given the large number of parameters we estimate, including occupation fixed effects in the wage regression. Since the LIAB does not contain information on AI adoption status, we restrict the sample to 2015–2019, prior to the diffusion of generative AI, to ensure that the estimated wage equation is not confounded by AI-induced wage changes. Worker characteristics $\tilde{X}_{it}$ include gender, education (three groups), age and age squared, experience and experience squared, and task-specific experience. Experience is constructed as the cumulative years of employment observed in the data. Exposure in task k 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.[12]

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,

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

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.[13]

4.2 AI Productivity

Measuring AI Productivity. We also construct a measure of AI productivity at the task level, Z_k^{AI}, 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).

To this end, we map task exposure into productivity using

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

[12] 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.

[13] Consistent with this assumption, in Appendix 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.

where e_k 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. 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.[14]

4.3 AI User Costs and Overhead Cost Shifters

We estimate the model at both the occupation level—fitting (6)—and the individual level—fitting (5)—to pin down AI task-level user costs and overhead cost shifters. For both approaches, we assume that overhead cost ε_i follows a truncated normal distribution with parameters (μ, σ), and that this 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).

The *location* of the deterministic AI gain and the overhead costs is only weakly identified, as the model is only informative about their relative levels through variation around the kink of the max operator in (5). In particular, when $\theta_k > \log(Z_k^{\text{worker}}/w_{io})$ holds for most k, i, separately identifying the common location of θ_k and the mean of ε_i becomes difficult.[15] Given this limited sensitivity, we fix the mean of ε_i at 5 for the reference worker, defined as male, low-educated, and in the youngest age group, while the mean overhead cost for other worker groups is shifted by $C(X_i)$ relative to this baseline; and verify post-estimation that our results are insensitive to this choice.

In terms of parameter *scale*, the externally constructed worker-side term $\log\left(Z_{ik}^{worker}/w_{io}\right)$ fixes the scale of the adoption gain. Hence the variance of the overhead cost is, in theory, identified.[16] However, the units of the externally constructed worker-productivity term carry

[14]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} outside the model serves to disentangle exposure effects from AI costs; but the precise parameterization of Z_k^{AI} is not critical for the estimation fit.

[15]To see this, note that in the "interior" case with $\theta_k > \log(Z_{ik}^{\text{worker}}/w_{io})$ *for all* k, i, adding a common constant c to all θ_k increases the deterministic gain by c, because $\sum_k \alpha_{ok} = 1$. Raising the reference mean of ε_i by the same c leaves the net adoption gain, and hence all adoption probabilities, unchanged.

[16]To see this, define $\theta_k \equiv \log\left(Z_k^{\text{AI}}/r_k\right), x_{ik} \equiv \log\left(Z_{ik}^{worker}/w_{io}\right)$, and write the deterministic gain from AI adoption as $G_i(\theta) \equiv \sum_k \alpha_{ok} \max\{\theta_k - x_{ik}, 0\}$. If the overhead cost is $\varepsilon_i = \mu + X_i'\gamma + \sigma u_i$, with u_i following

no independent meaning: any rescaling is observationally equivalent to a rescaling of the remaining parameters. In practice, we introduce and estimate a scale parameter ψ such that $\psi \cdot \log Z_{ik}^{\text{worker}}/w_{ik}$ is the worker productivity term. Once ψ is included, the model has a scale invariance.[17] Therefore the data cannot separately identify the scale of ψ, the scale of the task-level AI gains, and the dispersion of the overhead-cost shock. We must therefore choose one normalization and set the variance of overhead costs to $\sigma = 1$.[18]

Identification. We provide a formal identification argument for user cost, labor productivity scaling, and cost-shifter parameters $\Theta := (\{r_k\}_{k=1}^K, \psi, \gamma)$ based on observed occupation-level adoption shares. The idea is that under sufficient regularity on the distribution of ε_i and $Z_{ik}^{\text{worker}}/w_{io}$ (Assumption 1) and a rank condition on s (Assumption 2), occupation-level adoption shares $s(\Theta)$ can be locally inverted, recovering Θ. Our main result and intuition are stated here, while the details (including assumptions) are stated in Appendix D.

Proposition 1 (Local identification.)**.** *Suppose Assumptions 1 and 2 hold. Then, the parameter vector* $(\{r_k\}_{k=1}^K, \psi, \gamma)$ *is locally identified.*

The intuition is clear: The task-specific user costs $\{r_k\}_{k=1}^K$ are identified from cross-occupation variation in task composition (captured by variation in α_{ok}), conditional on worker composition X, AI productivity $\left\{Z_k^{AI}\right\}$, and worker productivity relative to wages $\left\{Z_{ik}^{\text{worker}}/w_{io}\right\}$. To see this, consider two occupations with the same worker composition and the same worker-side comparative-advantage terms, but different task weights: occupation A puts most weight on task 1, while occupation B puts most weight on task 2. If A exhibits a higher adoption share than B, then the model must assign a higher AI comparative advantage to task 1 than to task 2. Conditional on $\left\{Z_k^{AI}\right\}$, this means that task 1 must have a lower user cost r_1 than task 2. More generally, occupations that load more heavily on tasks with lower user costs should adopt more, and this cross-occupation covariance between adoption and task weights identifies the vector $\{r_k\}$.

a normalized distribution, then $P(D_i = 1) = F_u\left((G_i(\theta) - \mu - X_i'\gamma)/\sigma\right)$. Multiplying $(\theta_k, x_{ik}, \mu, \gamma)$ by $c > 0$ does *not* leave this adoption probability unchanged and so σ is identified.

[17] For any $a > 0$, AI adoption gains satisfy $G_i(a\theta, a\psi) = aG_i(\theta, \psi)$, so replacing $(\theta, \psi, \mu, \gamma, \sigma)$ by $(a\theta, a\psi, a\mu, a\gamma, a\sigma)$ leaves all adoption probabilities unchanged.

[18] Equivalently, one could impose $\psi = 1$ and estimate σ.

In turn, the parameter ψ is identified from cross-occupation variation in worker productivity relative to wages, holding fixed task composition, AI productivity and user costs, and overhead-cost shifters. The role of ψ is to determine how strongly worker-side comparative advantage matters for AI adoption. To see this, consider two occupations with the same task weights α_{ok}, the same AI productivity and user costs, and the same overhead-cost distribution, but different distributions of worker productivity relative to wages $\{Z_{ik}^{\text{worker}}/w_{io}\}$. Suppose workers in occupation A are systematically more productive relative to wages than workers in occupation B. Then AI has a weaker comparative advantage in occupation A, so A should exhibit a lower adoption share than B. A larger observed adoption gap between such occupations implies a larger value of ψ, because the model then needs worker-side productivity differences to have a stronger effect on the comparative-advantage terms $\theta_k - \psi \log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$. Conversely, if adoption rates are similar despite sizable differences in $Z_{ik}^{\text{worker}}/w_{io}$, the model implies a smaller value of ψ. Thus ψ is identified by how strongly cross-occupation differences in worker productivity relative to wages translate into differences in adoption shares.

Finally, the coefficients γ are identified from cross-occupation variation in worker composition, conditional on task composition and comparative advantage. Consider two occupations with the same task weights α_{ok}, the same AI productivity and user costs, and the same distribution of worker-side comparative advantage terms $\left\{Z_{ik}^{\text{worker}}/w_{io}\right\}$, but different education composition. If the occupation with more highly educated workers exhibits a higher adoption share, then this difference cannot be attributed to task-level profitability; it must be rationalized through the overhead-cost channel $C(X_i) = X_i'\gamma$. The model therefore assigns highly educated workers lower adoption frictions, i.e. a coefficient in γ that reduces overhead costs for that group. More generally, differences in worker composition across otherwise comparable occupations identify the coefficients in γ because they shift adoption through overhead costs while holding task-side profitability fixed.

Estimation. For the 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

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

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.[19]

The individual-level estimation is described in Appendix F.2.

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.

Nested Models. 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 $\{Z_k^{\text{AI}}\}_k$ and task importance weights $\{\alpha_{ok}\}_k$, which corresponds to occupational adoption rate (8). We next allow for 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).

We then incorporate heterogeneity in worker productivity $\{Z_{ik}^{\text{worker}}\}_{i,k}$, yielding the *Comparative Advantage* specification, which corresponds to occupational adoption rate (6). In this specification, workers with a comparative advantage in a given task—i.e., those who are

[19] 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.

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 and corresponds to adoption rate (6) plus cost shifters.

Model Selection. 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 producing an estimated parameter vector $\widehat{\Theta}$,

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

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 $\text{corr}\big(s_o^{\text{data}}, s_o(\widehat{\Theta})\big)$. To guide model selection among our nested models, following Andrews and Lu (2001), we compute

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

where $Q(\widehat{\Theta})$ 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; see Appendix F.2.

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 F.2 as robustness.

Model fit. Table 4 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 (column 1) and a correlation between actual and predicted occupational adoption shares of 0.56 (column 2). 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 only AI productivity that 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: the R^2 increases from 0.47 to 0.52, and the correlation performance measure gos up 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 4.

Appendix Table F.2 reports the results based on the individual-level estimation. The

Table 4: Model Fit: Occupation-level Estimation

Model	R^2	Corr.	AIC	BIC
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).

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 45-degree line would indicate a perfect model fit. The full model delivers a substantially better fit, consistent with Table 4. The improvement is most pronounced for occupations with extreme adoption shares—both at the low and high ends of the distribution.

Estimated parameters. We first report wage regression estimates of worker productivity, then turn to estimates of AI productivity, user costs, and overhead cost shifters.

Based on wage regression (9), 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, driven by heterogeneity both across and within occupations: The variance of mean log worker productivity across occupations is 0.05, while the average within-occupation variance is 0.08—indicating that dispersion within occupations exceeds variation across occupations.[20]

Turning to worker comparative advantage across tasks, worker productivity per unit wage,

[20]We also find substantial heterogeneity in occupational prices: The variance of the occupational fixed effect, FE_o, is 0.03, indicating an important role for occupational price differences in wage dispersion.

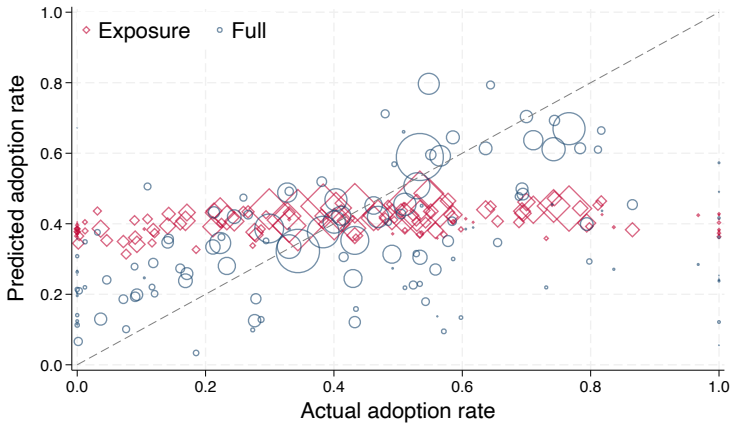


Figure 2: Model Fit for Different Model Specifications: Exposure vs. Full

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.

$\log\left(Z_k^{\text{worker}}/w_{io}\right)$, which we construct based on the wage regression estimates, varies substantially at the task level. The first column of Panel A in Table 5 reports, for each task, the mean of this ratio in logs—the human-side determinant of AI adoption. The gap between the highest-ranked task (Identifying Objects) and the lowest-ranked task (Manual Work Activities) is approximately 3.7 log points. Furthermore, this measure varies considerably across tasks within workers: the average within-worker variance in logs equals 1.98.

Next, Panel A in Table 5 reports estimates of task-level AI productivity, user costs, and their ratio under the *Full* model specification.

We find substantial heterogeneity in AI comparative advantage across tasks, indicated by a large variation in $\log(Z_k^{\text{AI}}/r_k)$. Cognitive tasks such as reasoning/decision (task 4) and technical activities (task 6) exhibit the highest productivity per unit cost, while manual work (task 5) and communication-intensive activities (task 7) 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.

Importantly, a stronger comparative advantage of AI does not necessarily arise from higher productivity alone. Tasks with high exposure-based productivity ranking do not always exhibit high productivity per unit cost once user costs are taken into account. For example, task 3 (Information Processing) has higher exposure than task 8 (Coordinating/Developing),

Table 5: Estimates of AI Productivity, User Costs and Worker Heterogeneity

Panel A: Task-level AI productivity and costs

	Tasks	$\log\left(\frac{Z_k^{\text{worker}}}{w_{io}}\right)$	Z_k^{AI}	r_k	$\log\left(\frac{Z_k^{\text{AI}}}{r_k}\right)$
1	Information Input	0.318	1.657	2.382	-0.363
				(2.660)	
2	Identifying Objects	2.229	1.314	0.040***	3.504
				(0.018)	
3	Information Processing	-0.210	2.112	0.029***	4.277
				(0.009)	
4	Reasoning/Decision	0.429	2.140	0.005***	6.150
				(0.002)	
5	Manual Work Activities	-1.467	1.038	0.095***	2.388
				(0.029)	
6	Technical Activities	-0.305	2.711	0.001***	8.388
				(0.000)	
7	Communicating and Interacting	-0.102	1.544	0.013***	4.809
				(0.004)	
8	Coordinating/Developing	-0.879	1.481	0.013***	4.725
				(0.004)	
9	Administering	0.144	1.767	0.009***	5.282
				(0.003)	

Panel B: Task-level worker heterogeneity and overhead-cost shifters.

	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 average log worker productivity per unit wage $\log\left(Z_k^{\text{worker}}/w_{io}\right)$, AI productivity ($Z_k^{\text{AI}}$), 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

indicated by $Z_3^{\text{AI}} > Z_8^{\text{AI}}$ in column 2, yet it 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, $r_3 > r_8$, in column 3.

The estimated user costs also vary substantially, and their ranking appears broadly consistent with features of the underlying task content. Technical Activities is the task with the lowest estimated cost, potentially reflecting that activities such as interacting with computers and documenting information are performed in digital environments where complementary infrastructure largely predates generative AI. A broad set of cognitive activities, including Reasoning/Decision, Administering, Communicating and Interacting, and Coordinating/Developing, lies in the middle of the distribution, suggesting a role for domain-specific customization. Information Processing has higher estimated costs, which may reflect the need to validate AI output, maintain accountability for final assessments as well as for evaluating quality and compliance. Identifying Objects and Manual Work Activities involve physical interaction and may require complementary capital investments, such as sensors or robotic actuators, driving up user costs further. Finally, the high estimated cost for Information Input is intuitive: many of these tasks involve monitoring, sensory observation, or equipment inspection, where AI use may require substantial verification and may be limited by privacy constraints.

Turning to worker heterogeneity in Panel B, 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.

In light of these estimates, we interpret the overhead costs as capturing worker-specific frictions in converting AI capability into productive use. These include learning costs, discretion to experiment, and the perceived career risk of relying on AI output and making mistakes. The estimated relationship with gender, age, and education should not be read as differences in technological ability, but as evidence that adoption frictions are shaped by workers' roles, autonomy, training opportunities, and organizational support.

Finally, we highlight the relative importance of worker heterogeneity across tasks for AI adoption, governed by the estimate of ψ. The standard deviation of the labor productivity per unit cost $\log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$ 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 Appendix Table F.1, we report summary statistics of labor productivity per unit cost across tasks.

Automation versus Augmentation. Our baseline model treats AI as replacing labor at the task level. However, whether AI primarily automates or augments labor is still an open question. To investigate this issue in our context, recall that we showed above that two alternative micro-foundations with labor-augmenting features generate a nested version of our baseline estimating equation, corresponding to the user-cost model in equation (7). The estimation results provide two pieces of evidence against the view that AI operates only through this purely labor-augmenting channel. First, the full model is selected over the user-cost model by our model-selection criteria. Second, in the full model we estimate $\psi > 0$, implying that variation in worker comparative advantage matters for adoption and that the nested user-cost equation is too restrictive.

Our interpretation is not that AI is purely labor-replacing, but rather that a purely labor-augmenting model is insufficient to rationalize the observed adoption patterns. While several micro-foundations—some under which AI automates and some under which AI augments labor—lead to the same reduced-form adoption equation as the user-cost model and therefore cannot be distinguished in theory based on information on AI adoption shares, the data favor a more general specification in which AI adoption depends on comparative advantage between AI and labor at the task level.

4.6 Robustness

Our baseline task classification relies on clustered GWAs, which are organized by *type of work* in O*NET. A natural concern is that this classification may not be the most informative one for measuring AI's comparative advantage over labor. In particular, if AI's comparative advantage were driven primarily by AI exposure rather than by type-of-work heterogeneity,

then a task classification constructed directly from AI exposure should provide a better fit to the data.

We assess this possibility by re-estimating the model using an alternative task classification based on *AI exposure*.[21] Specifically, we rank the 37 leaf GWAs by their mean AI exposure and partition them into nine equal-sized groups. Appendix Table F.5 compares the model fit under the baseline clustered-GWA classification and this exposure-based alternative. The exposure-based classification does not improve the fit on any metric, which implies that the type-of-work classification provides a useful basis for measuring AI's comparative advantage across tasks.

5 AI Adoption Driven by Comparative Advantage

We use our model estimates to shed light on three issues: how observed adoption patterns can be rationalized through our comparative-advantage framework; what exposure-based predictions miss relative to comparative-advantage-based predictions; and how future adoption is likely to evolve in the aggregate as well as across occupations and the wage distribution.

Our prediction exercises are deliberately narrow, both in terms of outcomes and time horizon. First, although we are comfortable using the partial-equilibrium model to generate predictions about AI adoption at a fixed output scale and prices, we do not use it to forecast labor demand. Second, we focus on short-run predictions. This choice reflects the fact that both the growth rates we feed into the model and the underlying technology are evolving rapidly, making longer-horizon projections considerably more speculative.

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 that have a high weight 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

[21] The two classifications differ substantially; see Appendix Table F.4.

Table 6, 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 6: 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 $\log\left(Z_k^{\text{AI}}/r_k\right) - \psi \log\left(Z_{ik}^{\text{worker}}/w_{io}\right)$, 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) high exposure but low adoption, (iii) low exposure and low adoption, and (iv) low exposure but

high adoption. Occupations in groups (i) and (iii) are well predicted by the exposure measure, while those in groups (ii) and (iv) represent systematic deviations driven by variation in AI user costs and worker productivity relative to labor costs across tasks.

The occupation "Office clerk" illustrates group (ii) with high exposure yet low adoption: It is intensive in the information-processing task, which exhibits high AI productivity but also high user costs, and therefore a modest comparative advantage of AI over labor. Consequently, adoption is lower 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, stemming from low user costs and low worker productivity, which results in greater adoption than exposure-based measures would suggest. Table F.6 reports the task weights for these occupations.

A substantial share of occupations falls into groups (ii) 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 (ii) 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.

Predicting AI Adoption. We use the estimated model to project the evolution of AI adoption rates under plausible trajectories of AI productivity and user costs over the next three years. To discipline the path of AI productivity, we use task-level growth rates derived from measures of AI *speedup*, defined as the ratio of human completion time without AI to that with AI assistance. These estimates are drawn from Appel et al. (2026), based on conversations with Claude.ai.[22] Using changes between different Claude versions from November 2025 (version 4) and February 2026 (version 5), we obtain monthly task-specific productivity

[22]This notion of time savings is consistent with the microfoundation of AI productivity in Appendix E.2.

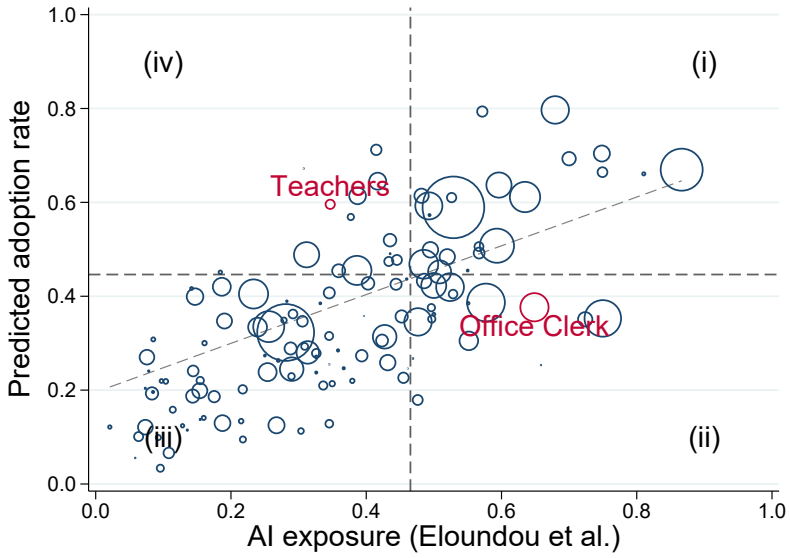


Figure 3: AI Exposure and Predicted Adoption

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.

growth rates, ranging from 3.62% for reasoning/decision-making to −3.81% for manual work activities; see Table F.7. We calibrate the path of user costs, r_k, to match observed aggregate adoption dynamics. Specifically, we target a monthly AI adoption growth rate of 2.27% in the US between November 2024 to November 2025, reported by the Real-Time Tracker (Bick et al., 2025). This implies an average monthly decline in user costs of 2.5%.

Figure 4 presents projected aggregate AI adoption rates over the three years following December 2024. Starting from the November 2024 adoption rate of 48%,[23] the model predicts aggregate adoption rising to 85% by end of 2027, driven by the joint evolution of AI productivity and user costs. Decomposing the contributions of each channel shows that both forces are quantitatively important, with reductions in user costs playing the larger role. Allowing only AI productivity to improve raises adoption to 65%, while allowing only user costs to decline raises adoption to 75%. This highlights the importance of user costs in addition to the capability of AI in shaping the diffusion of AI.

A central question in assessing the next stage of AI diffusion is: which occupations are

[23]We use the estimation results in Table 5 and the sample of workers in LIAB for the counterfactual exercises. Hence, the baseline adoption is slightly higher than 44% which is the adoption rate in DiWaBe due to differences in worker characteristics.

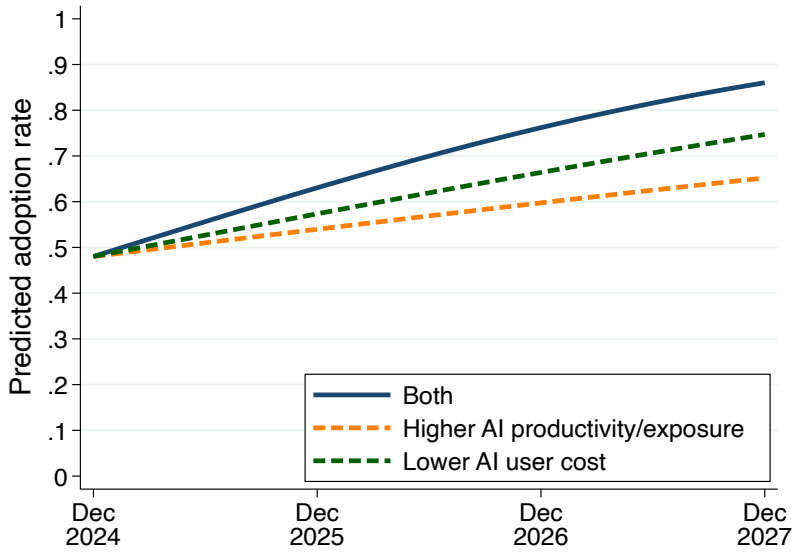


Figure 4: Projected AI Adoption Rate

Notes: The Figure plots projected AI adoption rates for three scenarios: both, Z_k^{AI} changes only, and r_k cost reduction only.

next? The left panel of Figure 5 examines cross-occupation heterogeneity by plotting model-predicted adoption rates three years ahead against current adoption rates. The model's answer to this question is not simply "those with the lowest current adoption," nor is it the occupations where adoption is already high. Instead, the largest predicted acceleration occurs among occupations in the middle of the current adoption distribution, such as office clearks or accountants. At the same time, occupations with high current adoption remain among those with the highest predicted adoption rates three years ahead, with IT occupations continuing to appear at the top of the distribution.

The right panel separates the roles of AI productivity and AI user costs by comparing three scenarios: AI productivity Z_k^{AI} change only (yellow), AI user cost r_k decline only (green), and both change (blue). The productivity channel generates increases in adoption that are heterogeneous across occupations depending on their task composition. For example, occupations in the middle of the current adoption distribution like accountants are among those with the biggest percentage-point increase in adoption, driven in large part by AI productivity growth. These occupations have relatively high shares α_{ok} in tasks such as identifying objects, information processing, reasoning/decision, which are among the task

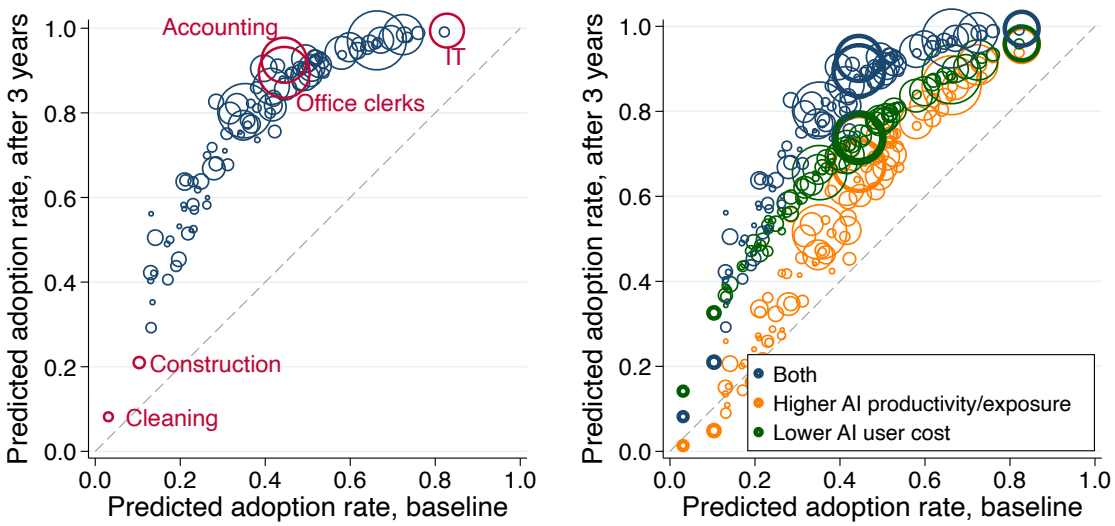


Figure 5: AI Adoption Rates across Occupations: Projection vs. Baseline

Notes: The left panel plots projected AI adoption rates against current AI adoption rates by occupation. The right panel reports the same relationship under three scenarios: both, Z_k^{AI} changes only, and r_k cost reduction only. Occupations highlighted in red on the left are displayed with bold outlines on the right.

categories with the fastest calibrated AI productivity growth. By contrast, occupations at the low end of the current adoption distribution such as construction workers and cleaners show essentially no gains through the productivity channel. These occupations rely heavily on manual work activities, for which our calibration implies little to no growth in AI productivity.

The user-cost channel generates a different pattern. Declines in AI user costs are a driving force behind increased adoption not only for occupations with currently moderate adoption but especially among occupations with lower current adoption rates. For these occupations, AI productivity is low today and is predicted to remain relatively low over the next three years, so a productivity-only forecast implies limited adoption growth. Once user costs fall, however, the model predicts sizable adoption gains. This distinction underscores why the next stage of AI diffusion depends on both improvements in AI capability and reductions in the costs of using AI.

AI adoption and inequality. Finally, we use the model to assess the distributional consequences of AI adoption. Specifically, we ask whether future diffusion is predicted to occur uniformly across the wage distribution or to remain concentrated among particular groups of workers. To do so, we examine the adoption–wage gradient using the binscatter in Figure 6.

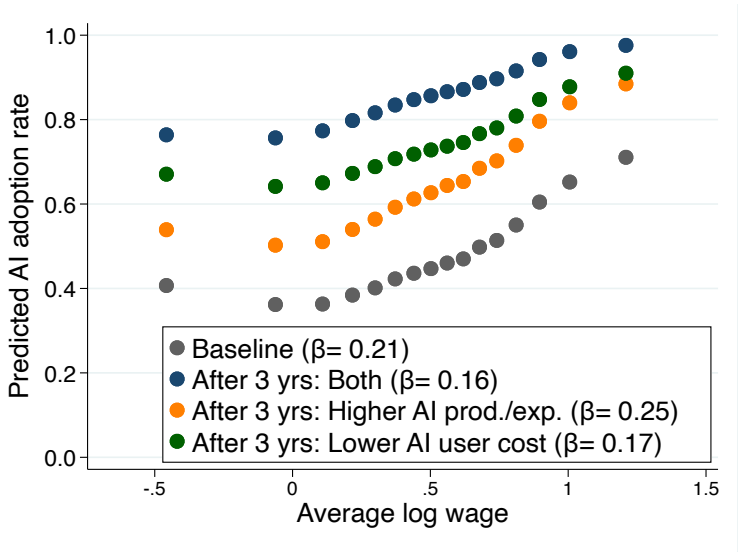


Figure 6: AI Adoption Rates across Wages

Notes: The figure is a binscatter of AI adoption rates against average log wages. It plots baseline adoption (grey) and projected adoption under three scenarios: both, Z_k^{AI} changes only, and r_k cost reduction only.

Baseline adoption, shown in grey, is tilted toward higher-wage workers: adoption rates increase with wages. The full-model prediction, shown in blue, indicates that AI diffusion is likely to become more broad-based in the near future, flattening the adoption-wage gradient.

The counterfactual predictions clarify the force behind this convergence. Comparing the exposure-only and user-cost-only counterfactuals shows that the flattening is driven primarily by falling user costs, shown in green, rather than by increasing AI exposure or productivity, shown in yellow. As Table F.7 shows, AI productivity grows fastest in tasks such as reasoning/decision, which are central to high-wage occupations. Thus, an exposure-only model would miss an important distributional implication of AI diffusion: adoption is likely to broaden across the wage distribution as declining user costs act as a force for convergence.

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 German worker survey 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.

Our forecasting exercise illustrates why the distinction between AI productivity and user costs matters for predicting AI diffusion. Projected adoption growth is not driven by improvements in AI productivity alone, nor by falling user costs alone. Rather, the largest adoption increases arise when AI becomes both more productive and cheaper or easier to use. This complementarity is especially visible across occupations: the occupations predicted to move next are primarily in the middle of the current adoption distribution, driven by both falling costs and rising productivity. In turn, at the lower end of the adoption distribution, falling user costs are predicted to generate sizable adoption gains in occupations in which AI productivity remains modest. Forecasts based only on productivity—or exposure—growth would therefore miss an important margin of near-term diffusion.

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.

Appel, R., M. Massenkoff, P. McCrory, M. McCain, R. Heller, T. Neylon, A. Tamkin, et al.: 2026, 'Anthropic Economic Index report: Economic primitives'. *Anthropic Economic Research*.

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.

Baslandze, S., Z. Edwards, J. Graham, T. McClure, B. H. Meyer, M. Sparks, S. R. Waddell, and D. Weitz: 2026, 'Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives'. Working Paper 34984, National Bureau of Economic Research.

Bick, A., A. Blandin, D. Deming, and N. Gazzaneo: 2025, 'Generative AI Adoption Tracker'. `https://www.genaiadoptiontracker.com/`. Harvard Project on Workforce. Data from the Real-Time Population Survey (RPS). Accessed: April 30, 2026.

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

Bick, A., A. Blandin, D. J. Deming, N. Fuchs-Schündeln, and J. Jessen: 2026b, 'Mind the Gap: AI Adoption in Europe and the US'. Technical report, National Bureau of Economic Research.

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

Bloom, N., J. M. Barrero, E. Brynjolfsson, S. J. Davis, B. Meyer, and E. Mihaylov: 2026, 'Firm Data on AI'. Working Paper w34836, National Bureau of Economic Research.

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: 2025, 'Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence'. *Digital Economy*.

Brynjolfsson, E., D. Li, and L. Raymond: 2025, '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.

Karger, E., O. Kuusela, J. Abaluck, K. A. Bryan, B. Halperin, T. R. Jones, C. Murphy, P. Trammell, M. Reynolds, D. Mayland, et al.: 2026, 'Forecasting the economic effects of AI'. Technical report, National Bureau of Economic Research.

Ledingham, D., T. Eloundou, S. Manning, P. Mishkin, and D. Rock: 2025, 'Beyond Automation: AI and Job Redesign'.

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.

Okamoto, M.: 1973, 'Distinctness of the eigenvalues of a quadratic form in a multivariate sample'. *The Annals of Statistics* pp. 763–765.

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} and we can define the labor unit cost

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

which equals the minimized task cost without AI.

When adopting AI, the task problem becomes

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

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

$$c_{k}^{AI} \equiv \frac{r_k}{Z_{k}^{\text{AI}}}.$$

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:

$$\mathcal{K}_{i}^{\text{AI}} \equiv \left\{k \in \mathcal{K} : c_{k}^{AI} < c_{ik}^{L}\right\} = \left\{k \in \mathcal{K} : \frac{Z_{ik}^{\text{worker}}}{w_{io}} < \frac{Z_{k}^{\text{AI}}}{r_k}\right\}.$$

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

Stage 2: Occupation-Level Cost Minimization. Let D_i be an AI adoption indicator. Given per-task unit cost, consider the occupation-level cost minimization problem for producing $y_{io} = y$ units for $d \in \{D_i = 0, D_i = 1\}$,

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

where the relevant unit costs are

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

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

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

The FOC for each task k is:

$$c_{ik}^{d} = \lambda \alpha_{ok} \frac{\prod_j y_{ij}^{\alpha_{oj}}}{y_{ik}} = \lambda \alpha_{ok} \frac{y}{y_{ik}}, \tag{11}$$

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

Rearranging (11) gives

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

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 task-level expenditures satisfy the share rule

$$c_{ik}^{d} y_{ik} = \alpha_{ok} \, C_{io}^{d}(y). \tag{12}$$

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 (12) 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,

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

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

$$UC_{io}^d = \prod_{k \in \mathcal{K}} \left(\frac{c_{ik}^d}{\alpha_{ok}} \right)^{\alpha_{ok}}, \tag{13}$$

which gives the unit costs described in the text.

A.2 AI Adoption Decision

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

$$\Delta_{io} \equiv \log UC_{io}^L - \log UC_{io}^{AI} \tag{14}$$

Using (13) and substituting the regime-specific unit cost,

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

Plugging in $c_{ik}^L = w_{io}/Z_{ik}^{worker}$ and $c_k^{AI} = r_k/Z_k^{\text{AI}}$ gives

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

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 k and the log of worker-i output per dollar in task k. 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 overhead cost $\exp(\varepsilon_i)$, 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 (14), the adoption indicator D_i is given by

$$D_i = \mathbf{1}\{\Delta_{io} > \varepsilon_i\},$$

and the adoption probability displayed in the text follows in a straightforward way.

B Alternative Models

In this section, we present three alternative models. The first imposes a labor constraint while keeping the production functions unchanged. The next two assume alternative production functions whereby AI augments labor at the task level rather than substituting for it. We show that all three models yield an estimating equation that closely follows (5); however, the interpretation of the estimates differs across the alternative micro-foundations.

B.1 Model with Labor Constraint

Most of the environment is identical to the baseline economy of Section 3; we therefore describe only the features in which this alternative model departs from it.

First, each worker is endowed with one unit of time, which is allocated across tasks subject to the constraint $\sum_k l_{ik} = 1$. Second, there is no fixed output scale. Third, instead of the stochastic multiplicative overhead cost, we assume that, prior to the adoption decision, worker i draws an idiosyncratic AI productivity term $\exp(-\varepsilon_i)$, where $\varepsilon_i \sim F_\varepsilon(\cdot \mid X_i)$. Conditional on adoption, the worker's output becomes

$$y_{io} = \exp(-\varepsilon_i) \cdot \prod_k (y_{ik})^{\alpha_{ok}}.$$

Finally, unlike the baseline model, this model requires no assumption on the structure of the labor market.

We next characterize the equilibrium.

No-AI benchmark. Absent AI, worker i in occupation o chooses the labor allocation across tasks,

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

where p_o denotes the output price. Given the Cobb–Douglas structure, workers allocate time proportionally to task weights, $l_{ik} = \alpha_{ok}$, so that the value of output equals

$$p_o y_{io} = p_o \prod_k \left(\alpha_{ok} Z_{ik}^{\text{worker}}\right)^{\alpha_{ok}}. \tag{15}$$

Worker's problem with AI. When adopting AI, worker i chooses both labor allocation and AI inputs,

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

where r_k denotes the unit cost of AI for task k. The associated Lagrangian, with multiplier λ on the time constraint, is

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

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

$$(\partial l_{ik}) \quad p_o \alpha_{ok} \frac{y_{io}}{l_{ik}} = \lambda \qquad (\partial K_{ik}) \quad p_o \alpha_{ok} \frac{y_{io}}{l_{ik} Z_{ik}^{\text{worker}}} Z_k^{\text{AI}} \leq r_k \qquad (k \notin \mathcal{K}_i^{\text{AI}}).$$

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

$$(\partial l_{ik}) \quad p_o \alpha_{ok} \frac{y_{io}}{Z_k^{\text{AI}} K_{ik}} Z_{ik}^{\text{worker}} \leq \lambda \qquad (\partial K_{ik}) \quad p_o \alpha_{ok} \frac{y_{io}}{K_{ik}} = r_k \qquad (k \in \mathcal{K}_i^{\text{AI}}).$$

AI adoption set. Combining the above conditions, for any $k_1 \in \mathcal{K}_i^{\text{AI}}$ and $k_2 \notin \mathcal{K}_i^{\text{AI}}$,

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

Reordering tasks by $\frac{1}{Z_{ik}^{\text{worker}}} \frac{Z_k^{\text{AI}}}{r_k}$ in decreasing order (so k_1 has the largest value), there exists a threshold k^* with $\mathcal{K}_i^{\text{AI}} = \{1, \ldots, k^*\}$. From the labor FOC, $\alpha_{ok} p_o y_{io} / l_{ik} = \lambda$, together with $\sum_k l_{ik} = 1$, the multiplier equals the shadow value of labor,

$$\lambda = (1 - \alpha_i^{\text{AI}})\, p_o y_{io}, \qquad \alpha_i^{\text{AI}} \equiv \sum_{k \in \mathcal{K}_i^{\text{AI}}} \alpha_{ok},$$

where α_i^{AI} is the total task share using AI. Substituting λ back yields the cutoff rule

$$\mathcal{K}_i^{\text{AI}} = \left\{ k \in \mathcal{K} \;:\; \frac{Z_{ik}^{\text{worker}}}{(1 - \alpha_i^{\text{AI}}) p_o y_{io}} < \frac{Z_k^{\text{AI}}}{r_k} \right\}. \tag{16}$$

Note that the adoption rule closely resembles (4); the only difference is that the market wage w_{io} is replaced by the endogenous shadow price of labor, $(1 - \alpha_i^{\text{AI}})\, p_o y_{io}$, which arises from the labor constraint. When the profits of all firms equal zero, and wages coincide with the shadow value of labor, $w_{io} = (1 - \alpha_i^{\text{AI}}) p_o y_{io}$, the adoption set is again (4).

Optimal inputs and equilibrium output. Conditional on $\mathcal{K}_i^{\text{AI}}$, optimal inputs are

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

Workers allocate labor across tasks according to task weights, but only among tasks that use labor; and demand more AI for tasks with larger weights α_{ok} and lower unit cost r_k. The

resulting output and the output value net of AI user cost are

$$y_{io} = \left(\exp(-\varepsilon_i) \prod_{k \notin \mathcal{K}_i^{\mathrm{AI}}} \left(Z_{ik}^{\mathrm{worker}} \frac{\alpha_{ok}}{1-\alpha_i^{\mathrm{AI}}} \right)^{\alpha_{ok}} \prod_{k \in \mathcal{K}_i^{\mathrm{AI}}} \left(\frac{\alpha_{ok} p_o Z_k^{\mathrm{AI}}}{r_k} \right)^{\alpha_{ok}} \right)^{\frac{1}{1-\alpha_i^{\mathrm{AI}}}}$$

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

AI adoption. Worker i adopts AI if and only if profits under adoption exceed the no-AI benchmark,

$$(p_o \exp(-\varepsilon_i))^{\frac{1}{1-\alpha_i^{\mathrm{AI}}}} \prod_{k \notin \mathcal{K}_i^{\mathrm{AI}}} \left(\alpha_{ok} Z_{ik}^{\mathrm{worker}} \right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\mathrm{AI}}}} \prod_{k \in \mathcal{K}_i^{\mathrm{AI}}} \left(\alpha_{ok} \frac{Z_k^{\mathrm{AI}}}{r_k} \right)^{\frac{\alpha_{ok}}{1-\alpha_i^{\mathrm{AI}}}} > p_o \prod_k \left(\alpha_{ok} Z_{ik}^{\mathrm{worker}} \right)^{\alpha_{ok}}. \tag{17}$$

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

$$\mathrm{P}(D_i = 1) = \mathrm{P}\left(\Delta_{io} > \varepsilon_i,\ \mathcal{K}_i^{\mathrm{AI}} \neq \emptyset \right) \tag{18}$$

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

and the occupation-level adoption rate is

$$s_o = \mathbb{E}[\mathrm{P}(D_i = 1) \mid o],$$

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

Reformulation of the adoption rule. The adoption condition above involves the set $\mathcal{K}_i^{\mathrm{AI}}$, which in principle requires solving a combinatorial problem over all subsets of tasks. However, this is not necessary in practice; 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 such that adoption condition (17) holds. Equivalently, this condition can be written as

$$\max_{\mathcal{K}_i \subset \mathcal{K}} \Delta_{io}(\mathcal{K}_i) > \varepsilon_i,$$

where $\Delta_{io}(\mathcal{K}_i)$ is analogously defined as above, but evaluated for an arbitrary subset of tasks $\mathcal{K}_i$, rather than the optimal AI adoption set $\mathcal{K}_i^{\text{AI}}$. Importantly, since $\Delta_{io}(\mathcal{K}_i)$ is additively separable across tasks in $\mathcal{K}_i$, the value-maximizing subset is the set of tasks whose term is positive, and in turn the adoption decision is equivalent to

$$\begin{aligned} \mathrm{P}(D_i = 1) &= \mathrm{P}\left(\sum_k \alpha_{ok} \max\left\{ \log \frac{Z_k^{\text{AI}}}{r_k} - \log \frac{Z_{ik}^{\text{worker}}}{\hat{w}_{io}}, 0 \right\} > \varepsilon_i \right) \\ \text{where} \quad \log \hat{w}_{io} &= \log p_o + \sum_{k'=1}^{K} \alpha_{ok'} \log(\alpha_{ok'} Z_{ik'}^{\text{worker}}). \end{aligned} \tag{19}$$

Here $\hat{w}_{io}$ is worker i's value marginal product of labor absent AI adoption. This representation shows that the adoption decision depends only on tasks for which AI has a comparative advantage over human, and it avoids explicitly solving for $\mathcal{K}_i^{\text{AI}}$; we therefore use this formulation for estimation. Notably, the resulting adoption probability (19) takes the same form as the one in our baseline model (5), with $\hat{w}_{io}$ replacing the wage w_{io}. While (19) is convenient for computing adoption probabilities, tasks for which the max operator is positive do not, in general, coincide with the optimal adoption set $\mathcal{K}_i^{\text{AI}}$; the representation follows from multiplying both sides of (17) by $(1 - \alpha_i^{\text{AI}})$, which preserves the adoption decision but not the composition of the optimal adoption set.

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 is

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

and thus simplifies to (7)—the user-cost model.

B.2 Model in Which AI is a Labor Productivity Improvement

Task-level output is produced by labor, $y_{ik} = Z_{ik}^{\text{worker}} l_{ik}$, prior to AI adoption, as before. However, we assume that AI is labor augmenting rather than labor replacing. Specifically, let AI adoption increase labor productivity from Z_{ik}^{worker} to $(1+\kappa_k)Z_{ik}^{\text{worker}}$, instead of fully replacing labor. Workers use AI for task k whenever $\kappa_k > 0$. We continue to assume that AI adoption comes with an idiosyncratic AI productivity $\exp(-\varepsilon_i)$, where $\varepsilon_i \sim F_\varepsilon(\cdot \mid X_i)$, which shifts total output, and focus on the problem of an individual worker with the same labor constraint as in subsection B.1. The optimal time allocation remains $l_{ik} = \alpha_{ok}$ as before. Substituting this allocation into profits under adoption and no adoption, worker i adopts AI if and only if profits are higher

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

Rearranging and taking the randomness of the productivity shock into account,

$$\mathrm{P}(D_i = 1) = \mathrm{P}\left(\sum_k \alpha_{ok} \max\{\log(1+\kappa_k), 0\} > \varepsilon_i\right).$$

With $\theta_k \equiv \log(1+\kappa_k)$, occupational adoption shares are given by

$$s_o = \mathrm{P}\left(\sum_k \alpha_{ok} \max\{\theta_k, 0\} > \varepsilon_i\right).$$

Thus, this adoption condition motivates the same empirical estimation specification as (7), which is itself nested by our baseline estimation equation (5). The interpretation of the estimates θ_k, however, differs from the baseline: rather than measuring the comparative advantage of AI over labor, here θ_k captures how much AI improves labor productivity.

B.3 Model with Cobb-Douglas Production for Task-Level Output

As in the baseline model, occupational output is produced by combining task-level outputs. But in contrast to the linear task production function from before, here task-level output is

produced using labor and capital based on a Cobb-Douglas technology:

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

where Z_k^K is the productivity of non-AI capital, and r_k^K is the unit price of non-AI capital. As in Appendix B.1, we consider the problem of an individual worker with labor constraint $\sum_k l_{ik} = 1$ and assume that the worker observes an idiosyncratic AI productivity shock $\exp(-\tilde{\varepsilon}_i)$ with $\tilde{\varepsilon}_i \sim F_{\tilde{\varepsilon}}(\cdot \mid X_i)$. This specification implies that labor $\sum_k l_{ik} = 1$ 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.

Without AI, the optimal non-AI capital demand and labor allocation are obtained by maximizing $p_o y_{io} - \sum_k r_k^K K_k^K$ subject to $\sum_k l_{ik} = 1$. Standard cost minimization for this nested Cobb-Douglas technology yields the output

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

Upon adopting AI, worker i uses AI for task k whenever 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

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

Worker i adopts AI when doing so raises the revenue net of capital costs, $p_o y_{io}^{\text{AI}} - \sum_k r_k K_k > p_o y_{io}^{\text{noAI}} - \sum_k r_k^K K_k^K$, which equals a constant fraction $1 - \xi$ of revenue under the Cobb-Douglas technology. Hence, the adoption condition reduces to $y_{io}^{\text{AI}} > y_{io}^{\text{noAI}}$.

$$\begin{aligned} \mathrm{P}(D_i = 1) &= \mathrm{P}\left(\Delta_{io} > \tilde{\varepsilon}_i, \right) \\ \Delta_{io} &\equiv \frac{\xi}{1-\xi} \sum_k \alpha_{ok} \max \left\{ \log \left(\frac{Z_k^{\text{AI}}}{r_k} \right) - \log \left(\frac{Z_k^K}{r_k^K} \right), 0 \right\}. \end{aligned} \tag{20}$$

Occupational adoption shares are given by

$$s_o = \mathrm{P}\left(\sum_k \alpha_{ok} \max\{\theta_k, 0\} > \varepsilon_i\right),$$

where $\theta_k = \log\left(\frac{Z_k^{\mathrm{AI}}}{r_k}\right) - \log\left(\frac{Z_k^K}{r_k^K}\right)$ and $\varepsilon_i = \frac{1-\xi}{\xi}\tilde{\varepsilon}_i$. As in the previous case, this condition leads to the estimating specification (7), a nested version of our baseline estimation equation (5). Here, however, θ_k measures the comparative advantage of AI over non-AI capital rather than over labor.

C Descriptive Evidence: Additional Details

C.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.

C.2 Additional Descriptive Evidence

Table C.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: The AI adoption share is the share of workers using AI at least sometimes for some task type. We report the 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 C.2: AI Adoption and Sources of Variation (Individual-Level)

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

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.

D Identification: Additional Details and Proof

Preliminaries. We formalize identification of the task-specific AI user costs $(r_1, \ldots, r_K)$, the labor productivity scale parameter ψ, and the overhead-cost shifter parameters $\gamma := (\gamma_1, \ldots, \gamma_J)$ in the occupation-level estimation based on equations (5) and (6). Throughout, we condition on objects constructed outside the model: the task weights $\{\alpha_{ok}\}_{o,k}$, AI productivity $\{Z_k^{\text{AI}}\}_k$, and the worker-side term $\log(Z_{ik}^{worker}/w_{io})$. The identification argument thus isolates the final-step recovery of r_k, ψ, and γ from revealed-preference adoption behavior.

For notational convenience, define

$$\theta_k \equiv \log\left(\frac{Z_k^{\text{AI}}}{r_k}\right), \qquad x_{ik} \equiv \log\left(\frac{Z_{ik}^{worker}}{w_{io}}\right).$$

Further, the overhead cost depends on observables through $C(X_i) = X_i'\gamma$, and its distribution is normalized by fixing the variance and the mean for the reference worker. Accordingly, write

$$\varepsilon_i = \mu_0 + X_i'\gamma + u_i, \qquad u_i \sim F_0,$$

where F_0 is the normalized cdf and μ_0 is fixed by normalization. Since Z_k^{AI} is known from the prior calibration step, there is a one-to-one mapping between r_k and θ_k, so local identification of r_k is equivalent to local identification of θ_k. We therefore work with $\theta := (\theta_1, \ldots, \theta_K)$ in place of $(r_1, \ldots, r_K)$ throughout. Letting $\Theta \equiv (\theta, \psi, \gamma)$ and

$$\Delta_i(\Theta) \equiv \sum_{k=1}^{K} \alpha_{o_i k} \max\{\theta_k - \psi x_{ik}, 0\} - \mu_0 - X_i'\gamma,$$

the occupation-level adoption share in (6) becomes

$$s_o(\Theta) = \mathbb{E}[F_0(\Delta_i(\Theta)) \mid o_i = o].$$

Identification Assumptions. We impose two identification assumptions. Denote the true parameter vector by $\Theta^0 = (\theta^0, \psi^0, \gamma^0)$.

Assumption 1 (Regularity)**.** *The cdf F_0 is continuously differentiable, and*

$$\mathbb{E}\left[\sup_{\Theta\in\mathcal{N}(\Theta^0)} |f_0(\Delta_i(\Theta))|\Big(1+\textstyle\sum_k |x_{ik}| + \sum_j |X_{ij}|\Big)\right] < \infty,$$

where $\mathcal{N}(\Theta^0)$ is a neighborhood of Θ^0. Moreover, for each occupation o, the conditional distribution of $(x_{ik})_k$ given o admits a continuous density, and $\psi^0 \neq 0$.

The second part of Assumption 1 rules out mass points at $\theta_k^0 = \psi^0 x_{ik}$, for all o and k. Together with the first part, this implies that the occupation adoption share $s_o(\Theta)$ is continuously differentiable in a neighborhood of Θ^0 despite the kink in $\max\{\theta_k - \psi x_{ik}, 0\}$. Hence,

$$\frac{\partial \Delta_i(\Theta)}{\partial \theta_k} = \alpha_{o_i k}\, \mathbf{1}\{\theta_k^0 > \psi^0 x_{ik}\} \quad a.s.,$$

so that differentiation passes through the expectation.

We further define the Jacobian blocks of $s_o(\Theta)$ with respect to θ, ψ and γ:

$$B_\theta(\Theta) \equiv \left[\frac{\partial s_o(\Theta)}{\partial \theta_k}\right]_{o,k}, \qquad B_\psi(\Theta) \equiv \left[\frac{\partial s_o(\Theta)}{\partial \psi}\right]_{o}, \qquad B_\gamma(\Theta) \equiv \left[\frac{\partial s_o(\Theta)}{\partial \gamma_j}\right]_{o,j}.$$

Computing all three blocks yields

$$\big[B_\theta(\Theta^0)\big]_{ok} = \alpha_{ok}\mathbb{E}\left[f_0\left(\Delta_i(\Theta^0)\right)\mathbf{1}\left\{\theta_k^0 > \psi^0 x_{ik}\right\} \mid o_i = o\right], \tag{21}$$

$$\big[B_\psi(\Theta^0)\big]_{o} = -\mathbb{E}\left[f_0\left(\Delta_i(\Theta^0)\right)\sum_k \alpha_{ok} x_{ik}\mathbf{1}\left\{\theta_k^0 > \psi^0 x_{ik}\right\} \,\Big|\, o_i = o\right], \tag{22}$$

$$\big[B_\gamma(\Theta^0)\big]_{oj} = -\mathbb{E}\left[f_0\left(\Delta_i(\Theta^0)\right)X_{ij} \mid o_i = o\right]. \tag{23}$$

We also define $A \equiv [\alpha_{ok}]$ and $M(\Theta^0) \equiv [m_{ok}(\Theta^0)]$, where

$$m_{ok}(\Theta^0) \equiv \mathbb{E}\big[f_0\big(\Delta_i(\Theta^0)\big)\,\mathbf{1}\big\{\theta_k^0 > \psi^0 x_{ik}\big\} \,\big|\, o_i = o\big],$$

so that $B_\theta(\Theta^0) = A \odot M(\Theta^0)$. Here, $\odot$ denotes elementwise multiplication.

We next aim to state an assumption under which the Jacobian of $s_o(\Theta)$ with respect to θ, ψ and γ has rank $K+1+J$. To do so we first establish the following lemma showing that,

as long as A has full column rank and $M(\Theta^0)$ has strictly positive entries, the first Jacobian block $B_\theta(\Theta^0)$ generically inherits the full column rank of A, i.e., $\text{rank}(B_\theta(\Theta^0)) = K$.

Lemma 1. *Let $A \in \mathbb{R}_+^{O\times K}$ with $O \geq K$ and $\text{rank}(A) = K$. The set*

$$\mathcal{Z}_A \equiv \left\{ M \in \mathbb{R}_{++}^{O\times K} : \text{rank}(A \odot M) < K \right\}$$

has Lebesgue measure zero in $\mathbb{R}^{OK}$. Equivalently, $\text{rank}(A \odot M) = K$ generically holds for $M \in \mathbb{R}_{++}^{O\times K}$.

Proof. Since $\text{rank}(A) = K$, there exists a set $\mathcal{O}_K = \{o_1, \ldots, o_K\}$ such that the corresponding $K \times K$ submatrix $A_{\mathcal{O}_K} \equiv (\alpha_{o_\ell k})_{\ell=1,\ldots,K;\, k=1,\ldots,K}$ is nonsingular: $\det(A_{\mathcal{O}_K}) \neq 0$. For any $M \in \mathbb{R}_{++}^{O\times K}$, write $M_{\mathcal{O}_K} \equiv (m_{o_\ell k})_{\ell,k=1,\ldots,K}$ for the corresponding submatrix, and define the determinant of the $K \times K$ submatrix of $A \odot M$ formed by the rows in $\mathcal{O}_K$,

$$p(M) \equiv \det(A_{\mathcal{O}_K} \odot M_{\mathcal{O}_K}).$$

We first show that p is a polynomial in the entries of M. By the Leibniz formula,

$$p(M) = \sum_{\sigma\in S_K} \text{sgn}(\sigma) \prod_{\ell=1}^{K} \left(\alpha_{o_\ell,\sigma(\ell)} m_{o_\ell,\sigma(\ell)}\right) = \sum_{\sigma\in S_K} \text{sgn}(\sigma) \left[\prod_{\ell=1}^{K} \alpha_{o_\ell,\sigma(\ell)}\right] \left[\prod_{\ell=1}^{K} m_{o_\ell,\sigma(\ell)}\right],$$

where S_K denotes the set of all permutations of $\{1, \ldots, K\}$. The bracketed coefficient $\prod_\ell \alpha_{o_\ell,\sigma(\ell)}$ is a constant, while each term $\prod_{\ell=1}^K m_{o_\ell,\sigma(\ell)}$ is a monomial in the entries of $M_{\mathcal{O}_K}$. Thus, p is a polynomial of degree K in the entries of M.

Next, we show that p is not identically zero. To do so, it suffices to find a matrix M at which p is nonzero. Indeed, take M to be the $O \times K$ matrix with all entries equal to one. Then $M_{\mathcal{O}_K} = \mathbf{1}_{K\times K}$, so $p(M) = \det(A_{\mathcal{O}_K} \odot \mathbf{1}_{K\times K}) = \det(A_{\mathcal{O}_K}) \neq 0$, since $\text{rank}(A) = K$.

By a standard result, the zero set of a nonzero polynomial on $\mathbb{R}^{O\times K}$ has Lebesgue measure zero (e.g., Okamoto, 1973). Therefore

$$\mathcal{Z}_p \equiv \left\{ M \in \mathbb{R}_{++}^{O\times K} : p(M) = 0 \right\}$$

has Lebesgue measure zero in $\mathbb{R}^{O\times K}$.

Finally, if $M \notin \mathcal{Z}_p$, then $A_{\mathcal{O}_K} \odot M_{\mathcal{O}_K}$ is nonsingular, so $A \odot M$ contains a nonsingular $K \times K$ submatrix and hence $\text{rank}(A \odot M) = K$. Therefore, $\mathcal{Z}_A \subseteq \mathcal{Z}_p$, and $\mathcal{Z}_A$ has Lebesgue measure zero. □

With $\mathcal{Z}_A$ defined above, we now state the local rank condition.

Assumption 2 (Local rank condition)**.**

(i) $A = [\alpha_{ok}]$ *has full column rank* K*;*

(ii) all elements of $M(\Theta^0)$ *are strictly positive, and* $M(\Theta^0) \notin \mathcal{Z}_A$*;*

(iii) $\text{rank}\big(M_{B_\theta}\left[B_\psi(\Theta^0)\ B_\gamma(\Theta^0)\right]\big) = 1+J$*, where* $M_{B_\theta} \equiv I_O - B_\theta(\Theta^0)\big(B_\theta(\Theta^0)'B_\theta(\Theta^0)\big)^{-1}B_\theta(\Theta^0)'$ *is the projection onto the orthogonal complement of the column space of* $B_\theta(\Theta^0)$*.*

Condition (i) requires that each task matters differentially across occupations. The strict-positivity part of condition (ii) imposes that the measure of workers near the adoption margin is strictly positive for every task in every occupation, which is naturally satisfied when the support of x_{ik} is sufficiently large. The remaining part of (ii) excludes the "Lebesgue-null" exceptional set identified by Lemma 1; this is a mild restriction, since the exceptional set has measure zero in $\mathbb{R}_{++}^{O\times K}$. Conditions (i) and (ii) together imply $\text{rank}(B_\theta(\Theta^0)) = K$, and combining this with condition (iii) yields the rank condition

$$\text{rank}\big(\left[B_\theta(\Theta^0)\ B_\psi(\Theta^0)\ B_\gamma(\Theta^0)\right]\big) = K + 1 + J. \tag{24}$$

We now discuss the primitives that ensure condition (iii). A condition for ψ to be separately identified is $\text{rank}\big(\left[B_\theta(\Theta^0)\ B_\psi(\Theta^0)\right]\big) = K + 1$. Comparing (22) with (21), $B_\theta(\Theta^0)$ depends on the weighted share of workers with $\psi^0 x_{ik} < \theta_k^0$, while $B_\psi(\Theta^0)$ further weights those workers by x_{ik}. Identification of ψ comes from cross-occupation differences in the productivity composition of marginal adopters.

Given the identification of θ and ψ, identification of γ requires that $B_\gamma(\Theta^0)$ contribute J linearly independent directions outside the column space of $[B_\theta(\Theta^0)\ \ B_\psi(\Theta^0)]$. From (23), the (o, j)-th entry of $B_\gamma(\Theta^0)$ is the weighted mean of X_{ij} in occupation o. The coefficient γ governs how worker characteristics X_i shift overhead costs and hence adoption probabilities.

Identification therefore requires that the composition of X_i vary across occupations in a way that shifts adoption through the overhead-cost channel and is not replicated by variation in task structure or worker productivity alone.

Proof of Proposition 1. Under Assumptions 1 and 2, the occupation adoption share $s(\Theta)$ is continuously differentiable in a neighborhood of Θ^0 and its Jacobian has generically full column rank $K + 1 + J$ at Θ^0. Local identification immediately follows from the Inverse Function Theorem. □

Remark. Here we translate identification of θ into identification of r. Because $\theta_k = \log Z_k^{\text{AI}} - \log r_k$, and Z_k^{AI} is treated as known, the mapping between θ_k and r_k is one-to-one. Hence local identification of θ is equivalent to local identification of r.

To express the Jacobian in terms of r, note that $\partial\theta_k/\partial r_k = -1/r_k$, so the Jacobian with respect to (r, ψ, γ) is

$$J_{(r,\psi,\gamma)} = \left[B_\theta(\theta,\psi,\gamma)\, D_r \;\; B_\psi(\theta,\psi,\gamma) \;\; B_\gamma(\theta,\psi,\gamma)\right], \quad \text{where } D_r \equiv \operatorname{diag}\left(-\tfrac{1}{r_1}, \ldots, -\tfrac{1}{r_K}\right).$$

Since D_r is nonsingular for $r_k > 0$, (24) implies that $J_{(r,\psi,\gamma)}$ also has full column rank $K+1+J$ at the true parameter.

E Estimation: Additional Details

E.1 Estimation of Worker Productivity

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

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

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

where K denotes the number of tasks; O the number of occupations; and $\tilde{J}$ the dimension of $\tilde{X}_i$.

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

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

For $j = 1$, $\tilde{X}_{it,j} = 1$, 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 2023. 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 in estimation that wages relevant for hiring decisions in 2023 have not yet been substantially affected by AI adoption and can be proxied by wages from 2015–2019.

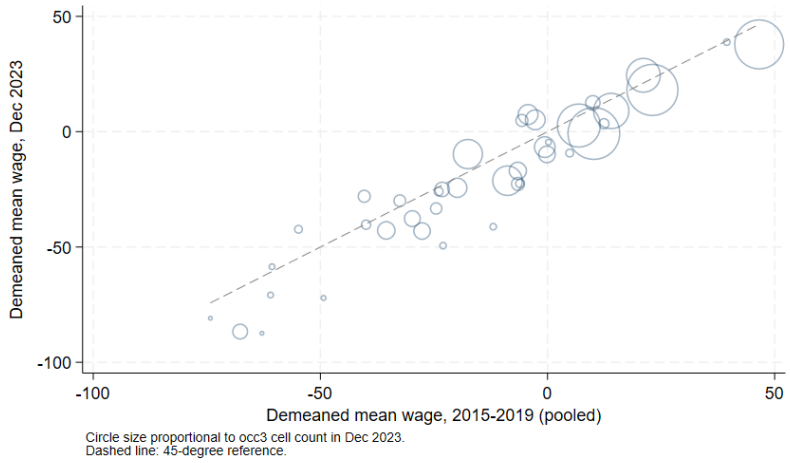


Figure E.1: Wages across Occupations: 2015–2019 vs. 2023

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

E.2 Estimation of 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

$$Z_k^{\text{AI}} = \frac{1}{\tau_k^{\text{AI}}} = \frac{Z_k^H}{1 - \Delta_k}. \tag{25}$$

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 $\Delta_k = e_k, \kappa = 1$, where e_k is the AI exposure measure developed by Eloundou et al. (2024) and reflects time savings from AI. 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. We use average values of e_k aggregated at the GWA level.

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 $\kappa < 1$ to reflect the average labor cost savings from AI in exposed tasks. For example, Acemoglu (2025) uses $\kappa = 0.27$, based on estimates from Noy and Zhang (2023) and Brynjolfsson et al. (2025).

Second, we can use survey-based measures of time savings, based on our DiWaBe data.

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 AI 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 the DiWaBe, the share of workers agreeing that AI has saved them time is $\bar{T} = 34.3\%$.

F Additional Estimation Results

F.1 Occupation-Level Estimation

Table F.1: 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.

F.2 Individual-Level Estimation

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

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

denote the model-predicted adoption probability; see (5).[24] 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:

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

Model Selection. 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

$$\mathrm{LR} = -2(\ell_0 - \ell_1),$$

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

The likelihood ratio test thus compares the fit of a restricted model to that of a more flexible unrestricted model that nests it. Because the unrestricted model can always weakly improve the maximized log-likelihood, the relevant question is whether this improvement is large relative to what would be expected from sampling variation alone under the restrictions. The test statistic allows us to assess whether the added parameters provide statistically

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

meaningful explanatory power.

We also compute the McFadden pseudo-R^2, defined as $R^2_{\text{McF}} = 1 - \frac{\ell(\widehat{\Theta})}{\ell_{\text{null}}}$, 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,

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

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

Results. Table F.2 reports the model fit of different specifications, and Table F.3 reports the estimation results.

Table F.2: Model Fit: Individual-level Estimation

Model	log-likelihood	LR p-val	McF R^2	AIC	BIC
Exposure	-4449	0.000	0.014	8900	8907
User Cost	-4348	0.000	0.036	8714	8776
Comparative Advantage	-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.3: Estimates of AI Comparative Advantage and Worker Heterogeneity

Panel A: Task-level AI productivity and costs

	Tasks	$\log\left(\frac{Z_k^{\text{worker}}}{w_{io}}\right)$	Z_k^{AI}	r_k	$\log\left(\frac{Z_k^{\text{AI}}}{r_k}\right)$
1	Information Input	0.318	1.657	0.149***	2.408
				(0.084)	
2	Identifying Objects	2.229	1.314	0.023***	4.040
				(0.014)	
3	Information Processing	-0.210	2.112	0.014***	5.051
				(0.006)	
4	Reasoning/Decision	0.429	2.140	0.002***	6.836
				(0.001)	
5	Manual Work Activities	-1.467	1.038	0.031***	3.498
				(0.011)	
6	Technical Activities	-0.305	2.711	0.004***	6.487
				(0.002)	
7	Communicating and Interacting	-0.102	1.544	0.021***	4.284
				(0.007)	
8	Coordinating/Developing	-0.879	1.481	0.010***	5.005
				(0.004)	
9	Administering	0.144	1.767	0.008***	5.437
				(0.004)	

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 average log worker productivity per unit wage $\log\left(Z_k^{\text{worker}}/w_{io}\right)$, AI productivity ($Z_k^{\text{AI}}$), 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F.3 Robustness

We rank the 37 leaf GWAs by mean exposure and partition them into nine equal-sized groups. Table F.4 reports all 37 leaf GWAs by exposure group, together with their baseline task group based on clustered GWAs. The two classifications differ substantially: for example, each exposure group draws from multiple baseline clusters. Yet, the fit is similar; see Table F.5.

Table F.4: Task Classification: Baseline vs. Exposure-Based

Group	Leaf GWA	Clustered GWA	AI Exposure
1	Performing General Physical Activities	Manual	0.008
	Operating Vehicles/Mechanized Devices	Manual	0.013
	Handling and Moving Objects	Manual	0.036
	Repairing/Maintaining Mechanical Equip.	Technical	0.053
2	Controlling Machines and Processes	Manual	0.108
	Coaching and Developing Others	Coord./Dev.	0.126
	Assisting and Caring for Others	Communicating	0.132
	Inspecting Equipment/Structures/Material	Ident. Obj.	0.193
3	Training and Teaching Others	Coord./Dev.	0.226
	Establishing/Maint. Interpers. Relat.	Communicating	0.267
	Monitor Processes/Materials/Surroundings	Info. Input	0.290
	Identifying Objects, Actions, Events	Ident. Obj.	0.335
4	Guiding/Directing/Motivating Subordinates	Coord./Dev.	0.356
	Selling or Influencing Others	Communicating	0.360
	Monitoring and Controlling Resources	Administering	0.360
	Estimating Quantifiable Characteristics	Ident. Obj.	0.365
5	Communicating w/ Supervisors/Peers/Subord.	Communicating	0.377
	Performing for/Working w/ the Public	Communicating	0.378
	Provide Consultation and Advice	Coord./Dev.	0.389
	Resolving Conflicts/Negotiating	Communicating	0.402
6	Staffing Organizational Units	Administering	0.406
	Scheduling Work and Activities	Reasoning/Dec.	0.415
	Organizing/Planning/Prioritizing Work	Reasoning/Dec.	0.432
	Judging Qualities of Things/Services/People	Info. Proc.	0.457
7	Getting Information	Info. Input	0.474
	Making Decisions and Solving Problems	Reasoning/Dec.	0.483
	Performing Administrative Activities	Administering	0.485
	Processing Information	Info. Proc.	0.509
8	Communicating w/ Persons Outside Org.	Communicating	0.513
	Analyzing Data or Information	Info. Proc.	0.516
	Updating and Using Relevant Knowledge	Reasoning/Dec.	0.516
	Evaluating Info. for Compliance	Info. Proc.	0.519
9	Interpreting Meaning of Information	Communicating	0.551
	Thinking Creatively	Reasoning/Dec.	0.556
	Developing Objectives and Strategies	Reasoning/Dec.	0.623
	Documenting/Recording Information	Technical	0.797
	Interacting With Computers	Technical	0.798

Notes: AI Exposure is the Eloundou et al. (2024) β rubric score, averaged across DWAs within each leaf GWA. Exposure-based groups rank all 37 leaf GWAs by exposure and divide them into nine approximately equal-sized groups (1 = lowest, 9 = highest).

Table F.5: Robustness I: Task Classification. Baseline vs. Exposure-Based

Specification	Q	RMSE	R^2	Corr$(s_o, \hat{s}_o)$
Clustered GWAs	3121	0.155	0.474	0.689
Exposure-based groups	3197	0.156	0.462	0.680

Notes: Both specifications use the user-cost model. *Clustered GWAs* partitions tasks by type of work; *Exposure-based* groups leaf GWAs by AI exposure rank; see Table F.4. $R^2 = 1 - \text{RMSE}^2/\text{Var}_w(s_o)$.

F.4 Additional Details on Post-Estimation Exercises in Section 5

Table F.6: Occupation Examples

Occupation	s_o	*Exp.*	$\hat{s}_o^{\text{full}}$	α_1	α_2	α_3	α_4	α_5	α_6	α_7	α_8	α_9
Office clerks, secretaries	38%	0.65	38%	0.06	0.02	0.20	0.12	0.04	0.15	0.20	0.09	0.13
Teachers	55%	0.35	60%	0.08	0.00	0.09	0.16	0.04	0.15	0.12	0.32	0.05
Average	45%	0.47	44%	0.08	0.07	0.11	0.15	0.08	0.13	0.16	0.16	0.06

Notes: s_o is the observed AI adoption rate. *Exp.* is the AI exposure index from Eloundou et al. (2024), which is between 0 and 1, with higher numbers indicating more exposure. $\hat{s}_o^{\text{full}}$ is the adoption rate predicted by our full model specification. α_k denotes the task weight for occupation o in task k: α_1 Information Input; α_2 Identifying Objects; α_3 Information Processing; α_4 Reasoning/Decision; α_5 Manual Work Activities; α_6 Technical Activities; α_7 Communicating and Interacting; α_8 Coordinating/Developing; α_9 Administering.

Table F.7: Task-Level AI Productivity Growth (Time-Saving)

k	Task	Time-saving (%)
1	Information Input	2.40
2	Identifying Objects	3.47
3	Information Processing	3.33
4	Reasoning/Decision	3.62
5	Manual Work Activities	-3.81
6	Technical Activities	1.87
7	Communicating and Interacting	1.06
8	Coordinating/Developing	0.50
9	Administering	1.85

Notes: The table reports the growth rates of *speedup* for each task, defined as the ratio of human completion time without AI to that with AI assistance, between November 2025 (version 4 of Claude.ai) and February 2026 (version 5 of Claude.ai); see Appel et al. (2026).