Pricing of Multidimensional Skills in IT Labor Market: Complementarities between AI and Traditional IT Skills

Shanjukta Nath* Jinsoo Yeo† Anandhi Bharadwaj‡ Sundar Bharadwaj§

March 4, 2026

Abstract

We quantify how the composition and complementarities of AI and IT skill bundles affect wages. Utilizing a transformer-based text analysis of job postings, we develop multivariate measures of skill intensity that capture implicit technological requirements often omitted in traditional keyword-based taxonomies. Our empirical strategy is grounded in a nested constant elasticity of substitution production function, which facilitates the estimation of heterogeneous substitution elasticities across three functional domains: IT Operations, IT Development, and AI. We find strong complementarities between AI and IT Development skills. Workers in the 90th percentile of IT Development skills gain a 18.5% wage premium from AI expertise, compared to just 5.1% for those with lower IT Development skills (10th percentile). In contrast, AI substitutes for routine IT Operations skills; however, operational expertise yields positive wage returns when both AI and Development skills are high. These complementarities are most pronounced among entry-level and mid-career workers, but diminish for experienced professionals. The results hold across work arrangements and locations, with the highest returns observed in tech hubs for in-person roles. Our findings highlight the importance of integrating AI and IT Development skills to maximize wage gains in the evolving IT labor market.

Keywords: AI skill complementarities, IT labor market, Wage premia, Compositional skill portfolios, IT Development skills, IT Operations skills, Transformer-based text analysis, Job posting data.

*Department of Agricultural and Applied Economics, University of Georgia (shanjukta.nath@uga.edu)

†Gary W. Rollins College of Business, University of Tennessee Chattanooga (jinsoo-yeo@utc.edu)

‡Goizueta Business School, Emory University (a.bharadwaj@emory.edu). Corresponding author.

§Terry College of Business, University of Georgia (sundar@uga.edu). All authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

The rise of Artificial Intelligence (AI) represents a significant shift in the economic landscape, with transformative implications for productivity and labor demand (Achiam et al., 2023; Bresnahan and Trajtenberg, 1995). Unlike the previous waves of automation, AI-enabled tools are influencing a broad spectrum of cognitive tasks, signaling a large-scale reconfiguration of labor rather than a simple reduction in headcount. This transformation is particularly salient in Information Technology (IT) occupations, which are intensive users of cognitive labor and among the earliest adopters of AI. While there is growing evidence of significant shifts in IT employment, the skill combinations through which AI reshapes the demand for different IT technical skill types and how these combinations translate into market wages remain open empirical questions.

The literature on skill-biased technological change (SBTC) documents that automation replaces routine tasks while augmenting complex, judgment-intensive activities, providing a foundation for examining how AI reshapes demand for IT skills (Acemoglu et al., 2023; Acemoglu and Restrepo, 2018; Dixon et al., 2021). However, aggregate occupational analyses often mask the nuanced pattern of labor market adjustment within job categories. For IT jobs, AI may simultaneously erode the market value of certain routine skills while increasing the premium for others (Acemoglu and Restrepo, 2018; Agrawal et al., 2019; Autor and Salomons, 2018; Bessen, 2019; Brynjolfsson et al., 2025; Eloundou et al., 2024). Prior research has been limited by the use of coarse, occupation-level measures or univariate AI exposure metrics, which fail to disentangle the economic returns to specific skill combinations (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024; Deming and Kahn, 2018). As a result, we know relatively little about how market wages reflect the evolving composition of skill portfolios within jobs.

We address this unresolved question by deploying transformer-based text analysis on a granular sample of IT job postings from 2018 to 2023. Our sample, drawn from the near-universe of job postings from Lightcast, is restricted to publicly traded US firms that report salary information, allowing us to examine wage outcomes and the marginal returns to multidimensional skill portfolios. Unlike earlier approaches that rely on predefined taxonomies, our methodology captures the joint structure of traditional IT capabilities and emerging AI expertise to construct multivariate, data-driven measures of skill demand. We develop a conceptual framework that models a firm's production function using a nested constant elasticity of substitution (CES) structure. The framework distinguishes three functional skill categories: IT Operations, IT Development, and AI.[1] These three skill measures are constructed by aggregating transformer-model-generated (Devlin, 2018; Vaswani et al., 2017) skill intensities across 30 granular technical skills, allowing substitution elasticities to vary across these groups, thereby accommodating heterogeneous complementarity and substitution among them. Linking technical skill requirements to their economic roles allows us to examine how AI interacts with different components of the IT skill bundle and how these

[1]Each of these skill categories is relevant for different tasks in the IT jobs. Operation skills focus on routine infrastructure and maintenance tasks, Development skills are for software, data engineering, and system design, and finally AI skills focus on advanced modeling and data-driven capabilities.

interactions translate into wage returns.

Our analysis reveals that wage dispersion in IT occupations is driven not by AI skills in isolation, but by their interaction with other components of the skill portfolio. We find strong evidence that AI complements IT Development skills while simultaneously substituting for more routine IT Operation skills. Specifically, workers in the 90th percentile of Development skills capture an 18.5% wage premium from AI expertise, whereas those at the 10th percentile gain only 5.1%. Conversely, AI weakens the valuation of standalone Operational skills. Notably, we document a positive three-way interaction, suggesting that operational expertise gains a wage premium only when embedded within broader skill bundles that include both Development and AI. Identification in the wage model leverages firm, state, and year-month fixed effects. The granular identification strategy is used to absorb unobserved firm heterogeneity, regional labor market dynamics, and aggregate time shocks, and the results remain robust across alternative specifications. Our strategy differs from most prior studies estimating the impact of skills on wages, which rely on between-firm, between or within-occupation or across MSA variation (Deming and Kahn, 2018; Deming and Noray, 2020; Hemelt et al., 2023; Josten et al., 2024; Woessmann, 2024).[2]

We further examine heterogeneity in AI-related wage premia across worker experience. The marginal effect of AI on wages is largest for junior roles across most skill portfolios, except the Operations-heavy portfolio, and remains positive but smaller for mid-career roles. The largest premia occur among Development-focused workers at both junior and mid-career levels (Brynjolfsson et al., 2025; Dell'Acqua et al., 2023). In contrast, the marginal effect diminishes for senior professionals, whose roles are less closely tied to specific technical skill bundles. Our analysis shows that these patterns align with a seniority-driven shift towards non-technical skills, such as management, cognitive, and social skills, which diminishes the AI-Development wage complementarity at senior career stages.

We also examine heterogeneity by work arrangement and job location. The marginal effect of AI on wages is largest for on-site positions, particularly within the pure IT Development portfolio when combined with AI skills. This complementarity between IT Development and AI also appears in remote and hybrid roles, though with smaller effects. The AI wage premium is further amplified for on-site roles located in tech hubs relative to those in non–tech hub areas.

From a managerial perspective, these findings suggest shifting talent strategies from occupational titles toward multidimensional skill portfolios. The results highlight the value of professionals who combine AI capabilities with system-level development and architectural expertise. While AI automates routine operational tasks, it increases the value of higher-order judgment and integrative problem-solving, particularly at early and mid-career stages. Accordingly, reskilling efforts should prioritize AI literacy alongside domain expertise that enables workers to operate effectively within this evolving task structure.

Overall, this study advances understanding of how AI reshapes labor demand within technical

[2] Woessmann (2024) provides a complete survey of papers on skills and wages, including the studies that use Lightcast job postings data.

workforces. Combining a nested CES framework with large-scale labor market data, we show how AI interacts with different segments of the technical IT skill bundle and how these interactions translate into wage outcomes. The results highlight that the labor market impact of AI depends on the composition of skill portfolios, not on occupational titles alone.

2 Related Work

Our review of the literature is organized around four thematic pillars at the intersection of technological change, skills, and labor markets. First, we examine the task-based framework of technological change to understand how automation reshapes the skill composition of labor demand. Second, we review the recent literature on AI adoption and its documented impacts on employment and wages. Third, we explore the literature on within-occupation skill heterogeneity, focusing on how bundles of social and technical skills drive wage premia. Finally, we discuss the methodological evolution of skill measurement, contrasting traditional dictionary-based approaches with modern transformer-based architectures.

The seminal literature (Acemoglu and Autor, 2011; Autor and Dorn, 2013; Autor et al., 2003) establishes that technology reshapes labor demand through task reallocation rather than simply eliminating occupations. Acemoglu and Restrepo (2019) and Acemoglu and Restrepo (2018) find that the displacement effects of automation are often countered by a reinstatement effect, where new task creation generates fresh demand for human labor. Extant research, operating at the level of occupations or task aggregates, offers limited insight into how specific skills or skill combinations are priced as technology evolves. They also fail to capture heterogeneity in returns across different skill combinations within the same occupation. In contrast, we demonstrate that wage effects are governed by skill portfolio composition rather than by occupational titles. By focusing within IT occupations, we are able to show that some technical skills appreciate in value while others depreciate simultaneously. This provides granular evidence of the task reallocation mechanism through the lens of skill pricing.

Recent empirical work documenting a dramatic increase in demand for AI skills finds that this surge is no longer confined to IT sectors but has diffused broadly into management, finance, and professional services. Alekseeva et al. (2021) document a substantial wage premium for job postings requiring AI skills, even after controlling for other high-level capabilities such as software proficiency. This suggests that while AI may automate certain high-skill tasks, the initial stage of its diffusion is creating a new tier of elite workers capable of managing and implementing these complex systems. Although this study provides micro-level evidence, it treats AI as a monolithic skill and does not consider how AI competencies integrate with or complement the existing portfolio of technical IT skills demanded by the labor market. Similarly, Tambe (2025) argues that AI enhances technological capabilities and calls for broad-based AI training across the firm. We add nuance to this finding by showing the asymmetric impact of AI skills, when combined with two distinct types of IT skills, on wage premia.

Acemoglu et al. (2022) documents the shift in AI postings over time, identifying an initial increase in 2010, followed by rapid growth starting around 2015–2016. Their findings indicate that AI exposure catalyzes a fundamental reconfiguration of task structures, driving a decline in legacy skills while creating new tasks that demand entirely new skill types. Although they document that AI-adopting establishments reduce their hiring of non-AI workers, suggesting a broad substitution effect, our research further investigates this trend. We show that this substitution is concentrated in IT Operations roles, whereas IT Development roles exhibit significant wage premia due to their complementarity with AI skills.

Expanding on the mechanisms of reallocation, Hampole et al. (2025) provides a granular assessment by distinguishing between "mean exposure," which drives substitution and reduces labor demand, and "dispersion in exposure," which allows workers to reallocate effort toward non-displaced tasks using the worker profiles data. Their findings indicate that while AI substitution is strong at the task level, aggregate employment effects are often muted because AI adoption generates significant firm-level productivity gains that increase overall labor demand. However, we find that these aggregate results can mask a critical divergence in skill valuations, with a bias toward complementary skills over substitutive skills. This results in significant wage premia only when AI is bundled with high-level architectural and development skills.

Finally, Babina et al. (2024) demonstrates that AI-investing firms experience fundamental shifts in workforce composition. While their work uses firm-level investment as a proxy, our study advances this by focusing on job-level skill bundle requirements. By decomposing AI effects by skill type, we move beyond the monolithic "AI shock" narrative to show that wage determination is a function of specific interactions between AI and traditional IT skills.

The growing emphasis on understanding how labor markets adjust to technological changes has also highlighted the rising importance of social and cognitive skills even as more advanced automation technologies have emerged. Deming (2017) shows that tasks that require tacit knowledge and interpersonal flexibility remain significant barriers to automation, leading the market to reward complementarity between social and cognitive abilities. Utilizing online job postings, Deming and Kahn (2018) find substantial heterogeneity even within narrowly defined roles, in which higher requirements for cognitive and social abilities are powerful predictors of both firm productivity and aggregate wage premia. Similarly, Hershbein and Kahn (2016) establishes that firms respond to economic shocks by changing their skill requirements, and these shifts directly impact wages. More recently, Mäkelä and Stephany (2024) finds a higher demand for social and cognitive skills such as resilience and analytical thinking in AI roles relative to non-AI roles. Although these studies establish the importance of examining specific skills and skill combinations, they primarily focus on broad, general skills (e.g., social, cognitive) rather than domain-specific technical skills associated with specific occupations. Our study complements the existing literature on skill heterogeneity within occupations by moving beyond broad skill categories to a more precise identification of 30 granular skills across the functional domains of AI, IT Development, and IT Operations. Our approach enables the generation of specific, testable predictions about the nature of complementarity

and substitution among these technical skill types. Because we are able to explicitly model higher-order interactions among the three skill types, we reveal critical skill synergies in which operational expertise, often viewed as a candidate for substitution by AI, gains economic value when bundled with both Development and AI capabilities.

The measurement of labor market demand for skills from unstructured job postings has traditionally relied on taxonomy-based approaches that match job posting text to predefined skill dictionaries. Data providers such as Lightcast (formerly Burning Glass Technologies) offer curated frameworks of standardized job titles and skill categories that are widely used in empirical research. Prior studies Alekseeva et al. (2021); Deming and Kahn (2018); Deming and Noray (2020); Hershbein and Kahn (2016); Mäkelä and Stephany (2024) typically use these Lightcast-generated skill labels of job postings (rather than the raw text) and map those to predefined skill dictionaries in order to quantify labor market demand for specific skills.

Recent research examining the labor market impact of artificial intelligence follows a similar approach. For example, Acemoglu et al. (2022) construct measures of AI exposure using Lightcast skill categories based on a finite set of explicit AI-related skill labels. Babina et al. (2024) measure AI overlap through the co-occurrence of granular skills associated with AI domains. Other studies employ TF–IDF methods or dictionary-based approaches to map job posting text to O*NET task descriptions.[3]

Although widely used, these approaches have several limitations in rapidly evolving technological domains. First, taxonomy-based systems rely on predefined skill categories that may lag behind emerging technologies. Capabilities such as "LLM fine-tuning" or "prompt engineering," for example, may be inconsistently represented or absent in existing taxonomies. Second, keyword-based methods rely on exact string matching and may therefore fail to capture implicit relationships among skills within job descriptions. Employers frequently omit foundational skills that are assumed given a particular technology stack. For instance, in our data nearly one quarter of job postings requiring "React Native" do not explicitly list "JavaScript," even though the latter is typically implied.

To address these limitations, we apply a zero-shot transformer-based classification framework using BART/BERT architectures to the raw text of IT job postings. A zero-shot approach is well-suited to this setting because it allows the classifier to recognize emerging skills without requiring task-specific labeled training data. Rather than relying solely on predefined skill labels, the transformer model uses contextual language representations (Vaswani et al. (2017)'s attention mechanism in the transformer framework) to infer underlying skill requirements and assigns probability scores that reflect the intensity of different skills, even when specific labels are not explicitly present in the text. We aggregate these granular skill intensities (zero-shot probabilities) into three functional domains, AI, IT Development, and IT Operations, allowing us to construct multidimensional measures of IT skill portfolios. This structure enables a compositional analysis of IT

[3]TF-IDF methods stand for Term Frequency-Inverse Document Frequency, and these are pre-transformer methods for text analysis.

skill portfolios and allows us to study complementarities across AI, Development, and Operations skills. Because the framework does not require task-specific training, it can scale to large volumes of job-posting data while maintaining consistent classification performance across emerging skill categories.

3 Data

We draw on two datasets for this study. The primary job posting dataset is from Lightcast, widely used in the labor economics literature as a near-universe of online job vacancies (Deming and Kahn, 2018). Relative to occupational data sources such as O*NET, which relies on expert surveys and infrequent periodic updates, posting-level data provides more timely and granular insights into evolving labor demand. Lightcast provides particularly strong coverage of the IT labor market, with IT postings captured more consistently than sectors such as agriculture and blue-collar occupations (Hershbein and Kahn, 2018). Our analysis is further restricted to postings with explicitly reported salary information, as wage information is central to our study. This is an important distinction, as other labor market data providers often rely on imputed salary measures, which can introduce substantial measurement error (Dai et al., 2025).

Lightcast provides a comprehensive set of cleaned and structured posting-level variables, including posting dates, standardized job titles, employer identities, skill labels, and detailed geographic regions. While extant literature has primarily focused on this structured information to study AI adoption and labor market dynamics (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024), our empirical strategy focuses on the "raw text" of the job postings to construct multivariate latent IT skill measures. We complement the Lightcast data with firm-level financial characteristics from Compustat, which is matched to job postings data using company names.

Sample construction Our initial dataset comprises more than 234 million job postings spanning 72 months, from January 2018 to December 2023. From this universe, we construct a focused sample of 870,810 IT job postings from 2,709 unique U.S. publicly traded firms that report salary information, with IT occupations identified using the two-digit O*NET code 15. To link job postings to firm-level financial data, we implement a hybrid matching procedure integrating semantic vector search, LLM-assisted review, and human validation. Specifically, we generate dense text embeddings for company names using a BERT-family sentence transformer (all-MiniLM-L6-V2), indexed with Facebook AI Similarity Search (FAISS) to identify the nearest Compustat candidates based on cosine similarity. Relative to fuzzy string matching approaches commonly used to link job postings data to Compustat (Babina et al., 2024; Fedyk and Hodson, 2023), this embedding-based approach enables semantic matching, capturing corporate aliases and parent-subsidiary relationships with limited lexical similarity. Candidate mappings are then reviewed by GPT-4o, anchored to August 2024 to ensure consistency across corporate structures and M&A activity, with low-confidence matches subject to final human expert validation. This hybrid AI-expert procedure

is consistent with emerging standards for maintaining large-scale occupational databases (ONET Development, 2025). Table A1 presents the mapping results and summarizes the sample by year.

3.1 Transformer-Based Multi-dimensional Skills for IT Jobs

Given the rapid changes in the demand for skills in the current labor market, mapping job posting text data to different skills poses a critical challenge for researchers. To address this challenge, we introduce a new measure of AI and traditional skills in IT jobs using a zero-shot text classification (Yin et al., 2019) method. The classifier uses a transformer-based framework pretrained on natural language inference tasks to map unstructured job posting text to user-defined skill labels. The method has the following strengths.

First, unlike keyword-matching or dictionary methods, the transformer's attention mechanism leverages bidirectional context to capture underlying semantic connections. This allows our model to assign probability scores for skill intensity even when labels are not explicitly present in the job-posting text. Second, by utilizing a data-driven approach to induce granular IT skill categories, which are further aggregated to three functional domains (AI, IT Development, and IT Operations), our method enables a compositional analysis of the overall skill landscape of IT jobs. Without requiring any task-specific training, our approach enables the automated processing of high volume postings data while maintaining high precision.

The most significant advantage of this approach is that the model is pre-trained and does not need to be trained on the job postings data to develop the classifier. This, in turn, means the researcher does not need to perform the costly task of relabeling and retraining the model as rapidly changing skills emerge in the current labor market transformation following the introduction of Generative AI. Although job postings capture hiring intentions rather than realized worker tasks, they provide a widely used revealed-preference measure of the skills firms demand in the labor market and therefore offer a direct signal of employer demand for specific skill combinations.

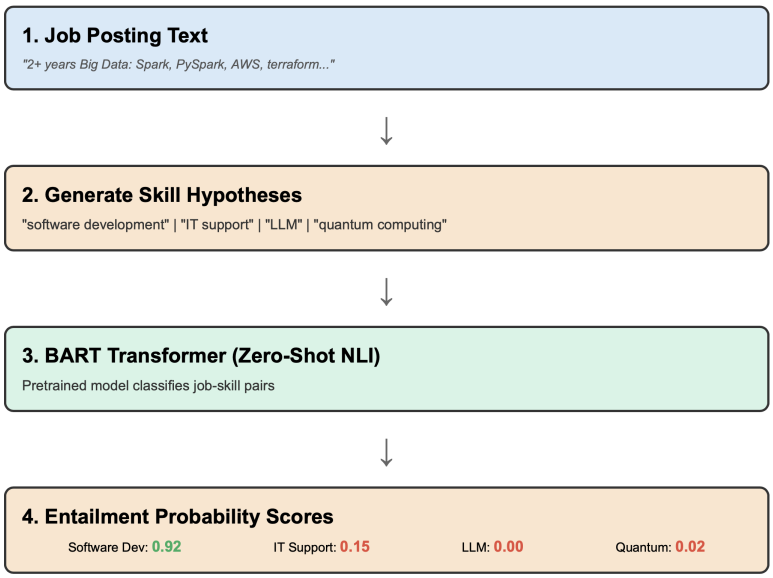
We explain the steps for zero-shot classification using an example from our data-set. Consider the following abbreviated version of a job text "*BS in Computer Science or Engineering with at least 2 years of professional software work experience. 2+ years of experience in Big Data Platform technologies required Spark or PySpark or Scala. Experience with tools such as Databrick or DeltaLake preferred. Working knowledge of Amazon cloud technologies. Strong collaboration and teamwork skills...*" and to keep the illustration simple, we work with four skill labels: Software development, IT support, LLMs, and Quantum computing. [4]

In the first step, the method creates four hypotheses using the text and the labels: i) This job requires Software development skills, ii) This job requires IT support skills, iii) This job requires LLM skills, and iv) This job requires Quantum computing skills. Next, it formats the text input (premise) and the hypothesis into a premise-hypothesis pair, which can be used as an input

[4] In our empirical application, we work with 30 skill labels encompassing AI and traditional skills in IT jobs for publicly listed US firms. This job text is an abbreviated version of a longer body of a job posting from the Lightcast dataset.

into the transformer model. In the third step, a pretrained transformer model such as Facebook BART/BERT will use the attention mechanism (Vaswani et al., 2017) to extract the underlying semantic connections between the text and the labels. If there is a strong semantic connection between a label and the text, it will be assigned a higher probability score. Figure 1 summarizes each step for the classification pipeline. The classification was performed on a university high-performance computing cluster, and processing the full data required approximately 375 GPU hours in total, with each month's job posting classification taking approximately six hours.

Figure 1: Zero-shot classification pipeline



Our classification method differs from related work that has used job postings and resumes to estimate skill demand in several ways. First, the attention mechanism in Vaswani et al. (2017) and the transformer framework Devlin (2018) do not use keyword-matching/dictionary methods that search for exact strings of skill labels in job-posting text. Related papers that use job-posting data use the skill categories provided by Lightcast for each job posting and construct AI and other social and cognitive skill measures using the Lightcast-generated skill taxonomies (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024; Deming, 2017; Deming and Noray, 2020).

Keyword matching techniques or dictionary methods completely ignore the context of the text, and hence, they cannot learn about the underlying semantic relationships between text and labels. The attention mechanism uses bidirectional context (left and right) to capture semantic connections in the text. Even when skill labels are not explicitly present in the job text, the attention mechanism can capture the text's underlying semantic connections to the labels. This context-awareness is essential for measuring the demand for modern heterogeneous IT labor skills, in which skill expression is often implicit, rapidly evolving, and loosely standardized.

As the attention mechanism leverages semantic depth rather than literal string matching, the selection of input labels constitutes a pivotal design choice in our empirical framework. These labels

serve as the semantic anchors for our zero-shot classification algorithm, allowing us to delineate the boundaries of modern IT labor. Given the inherently fluid nature of these boundaries, which legacy taxonomies often fail to capture in detail, we adopt a systematic, data-driven approach to label construction. Specifically, we derive a set of 30 IT skill categories using principled, iterative prompting from a Large Language Model (LLM).

Existing vendor-provided taxonomies highlight this challenge. For example, Lightcast currently characterizes the IT sector using 69 detailed categories, each intended to represent a distinct IT specialization. Yet even such fine-grained classifications are not stable. As an illustration, Computer Vision, a core AI domain and a central component of AI measure in Babina et al. (2024), is classified under an "Image Analysis" subcategory, positioned far from AI and even from IT more broadly. Such inconsistencies underscore the limitations of static taxonomies for capturing contemporary IT skill demand.

To construct a structured representation of the latent multivariate IT technical skill domain, we utilize an occupational-specialist prompting framework with LLM. Specifically, the LLM is instructed to assume the role of an IT occupational expert and to propose distinct skill categories for traditional IT and emerging AI domains. For each category, the model must provide: (1) a formal domain name, (2) a brief rationale for its relevance in labor market analysis, and (3) a list of representative software tools or technical infrastructures associated with that domain. We choose to derive 25 AI-specific and 25 non-AI-specific IT skill categories separately. This distinction is necessitated by the role of AI as a general-purpose technology (Goldfarb et al., 2023), where AI capabilities increasingly permeate traditional IT functions and the specific nuances of emerging AI specializations are subsumed by broader legacy IT categories.[5] To assess the robustness of these categories, we conducted sensitivity testing across multiple LLM architectures and prompt variations.

Our goal is to ensure a framework that approximates representative and collectively exhaustive skill domains, without relying on predefined vendor taxonomies. This requirement is essential for robustly quantifying the heterogeneity of IT labor and capturing the latent complementarities between existing IT expertise and AI capabilities. To support stability, we conduct large-scale iterative testing by running the induction prompt 1,000 times.[6]

Each induced category is accompanied by a technical justification and subsequently subjected to a Human-in-the-Loop (HITL) cross-validation process. Using the resulting frequency distributions, we iteratively refine the taxonomy by removing infrequently appearing categories and excluding overly granular labels that correspond to specific tools or individual skills rather than broader skill categories. Analogous to the elbow method in cluster analysis, we evaluate the trade-off between

[5]We further impose negative constraints, explicitly excluding non-technical "soft skills" and restricting attention to technical IT skill categories and frontier AI domains.

[6]While sampling randomness is theoretically suppressed under greedy decoding ($T = 0$), residual stochasticity may arise in LLM outputs, precluding strict deterministic results. We address this variability through frequency-based refinement and semantic consistency checks, ensuring that our final taxonomy represents a stable, LLM consensus-driven framework.

category granularity and semantic distinctness to arrive at a final set of 30 skill domains.[7] This frequency-guided pruning process yields a parsimonious taxonomy that balances semantic precision with empirical coverage, ensuring that our measures of IT skill demand are both empirically robust and stable across the near-universe of observed job postings (Figure A1).[8]

3.2 Summary Statistics

This section reports summary statistics for the key variables used in our analysis, drawn from the merged Lightcast-Compustat dataset. The statistics describe the distribution of job-posting-level outcomes, firm characteristics, and the zero-shot-generated skill measures that form the basis of our empirical exercises. All job postings are restricted to publicly traded firms with valid Compustat identifiers, ensuring consistent linkage between labor demand and firm-level information. Summary statistics are reported at the job-posting level unless otherwise noted.

To identify the latent structure of IT labor demand and map technical requirements to economic functions, we aggregate the granular skill measures into three functional categories, IT Operations, IT Development, and AI skills, corresponding to infrastructure and maintenance tasks, software engineering and system design, and advanced modeling and data-driven capabilities respectively. These three functional categories constitute the core explanatory variables in our analysis, and salary information is the core dependent variable.

IT Operations (O), include skills related to maintenance and operations of systems. Precisely, this includes IT Support, Systems Administration, Network Administration, Networking, Database Management, Operating Systems, Virtualization, and Cybersecurity. These skills will enable the performance of routine IT tasks. Second is IT Development (D), which involves skills in digital engineering and software development and requires more creative and complex problem-solving. This includes Software Development, Web Development, Mobile Development, Mobile App Development, Cloud Computing, DevOps, Big Data, Data Engineering, Embedded Systems, and Internet of Things (IoT). Finally, all the AI skills are grouped into AI, including Artificial Intelligence, Machine Learning, Reinforcement Learning, Computer Vision, Quantum Computing, Robotics, Blockchain, Blockchain Technology, Generative AI, Natural Language Processing, Large Language Models, and Data Science.

The longitudinal data on job postings reveal a stark decoupling of skill demands within the IT sector. As shown in Figure 2, aggregate AI, Development, and Operations skill intensities follow distinct trajectories throughout the sample period. While AI intensity rises steadily before reaching a high-level plateau, demand for traditional IT Development and IT Operations declines. This

[7] The 30 skill categories include: Artificial Intelligence, Big Data, Blockchain, Blockchain Technology, Cloud Computing, Computer Vision, Cybersecurity, Data Engineering, Data Science, Database Management, DevOps, Embedded Systems, Generative AI, Internet of Things (IoT), IT Support, Large Language Models, Machine Learning, Mobile App Development, Mobile Development, Natural Language Processing, Network Administration, Networking, Operating Systems, Quantum Computing, Reinforcement Learning, Robotics, Software Development, Systems Administration, Virtualization, and Web Development.

[8] The current taxonomy represents the baseline induction from our prompting framework. While certain categories exhibit semantic proximity, they represent distinct nuances in the model's consensus induction. We remain open to consolidating these into broader clusters for further sensitivity testing in future iterations.

divergence emerges prior to late 2022 and persists thereafter, underscoring substantial heterogeneity beneath aggregate trends. The portfolio-level panels further show that skill demand is dynamically recomposing within IT occupations, with sharp differences in timing and magnitude across specific domains rather than a uniform expansion or contraction of IT demand. Consistent with these patterns, Figure A2 documents a concurrent shift in the composition of IT job postings toward higher salary quartiles, suggesting a reallocation of labor demand toward higher-paying IT roles.

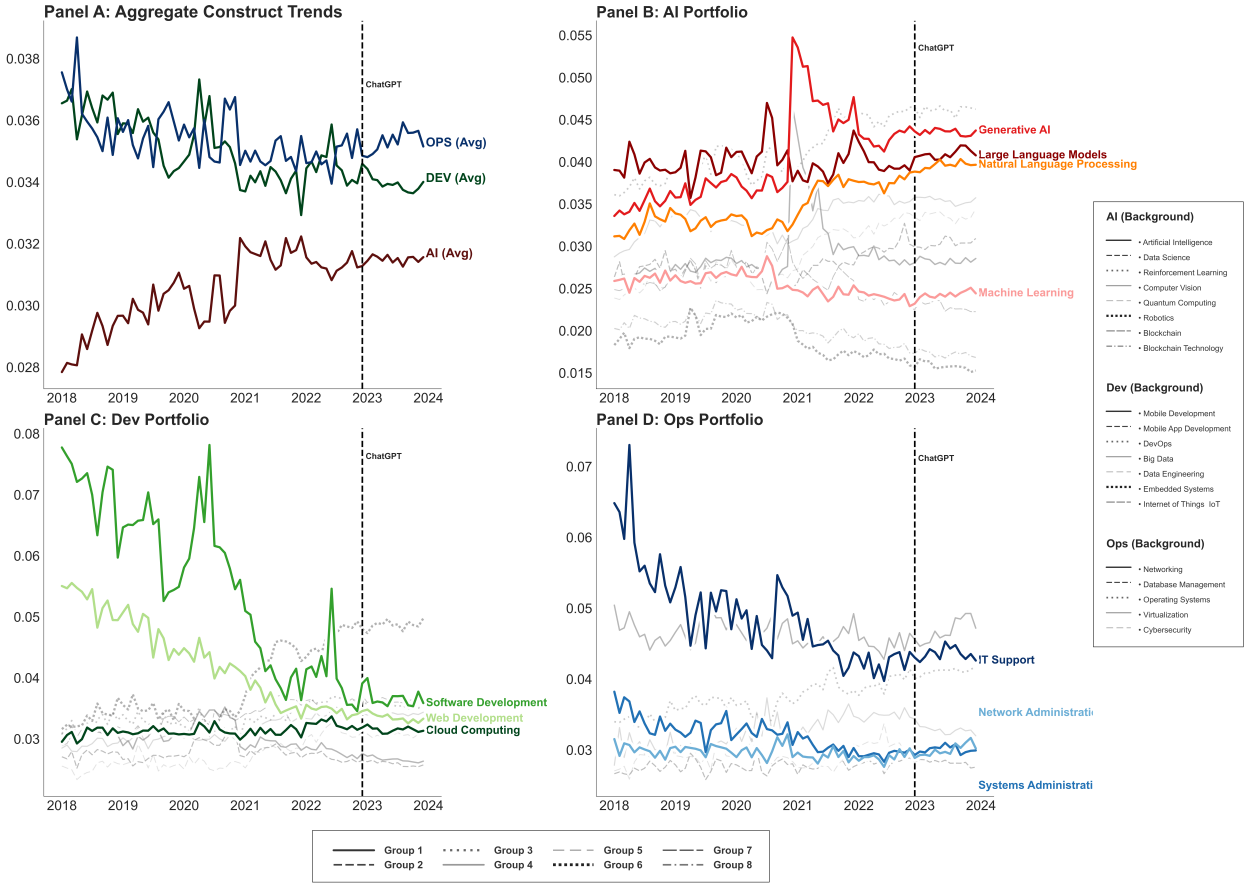
Collectively, these figures highlight substantial heterogeneity in the evolution of skill demand across the full set of granular, zero-shot-generated skill measures. While this high-dimensional representation captures rich variation in skill demand, it complicates interpretation and obscures recognition of broader economic patterns. Hence, we conduct our empirical analysis using the three interpretable aggregates of skill demand.

A potential concern using salary data is the non-random nature of salary disclosure. Our primary analysis sample is necessarily smaller than the full universe of postings, as it restricts attention to publicly traded firms, IT occupations, and postings with usable salary information. To assess whether this selection induces systematic bias, we conduct a comprehensive representativeness audit comparing the analysis sample (n=870,810) with the full universe of IT job postings (n=5,208,603). Following standard econometric diagnostics, we compute standardized mean differences (SMD) across key job-level and firm-level characteristics (see Table A2).

The balance results indicate that salary-disclosing postings are highly comparable to non-disclosing postings along core dimensions. Specifically, the SMDs for minimum experience requirements (0.02), educational thresholds (0.06), and skill count (0.03) are all well below the conservative 0.1 threshold typically used to denote meaningful imbalance. While salary disclosure is more frequent among fully remote roles and within firms of modestly smaller scale (likely reflecting strategic transparency by agile firms competing for national talent) the underlying human-capital requirements remain invariant. To ensure these modest differences do not confound our results, our primary specifications in Section 6 include stringent controls for unobserved firm differences (by incorporating firm fixed effects), remote-work status, and geographic-transparency mandates (state or MSA fixed effects). Taken together, these results confirm that our analysis sample is highly representative of the broader IT labor market, alleviating concerns that the estimated wage premia are driven by selective disclosure or unobserved job-level heterogeneity.

Table A3 reports summary statistics for the main analysis sample. Panel A presents the distribution of posted salaries, which serve as the dependent variable in the wage regressions. The statistics indicate substantial variation in wages across job postings. Panel B summarizes the zero-shot-generated skill scores aggregated into three categories: IT Operations, IT Development, and AI skills, which constitute the primary explanatory variables of interest. The reported means and dispersion show meaningful variation in skill intensity across postings. Panels C through E report summary statistics for education requirements, experience thresholds, and work arrangements, which are included as control variables in the baseline specifications and then used for heterogeneity analysis in section 7.

Figure 2: Temporal Trends in IT Skill Portfolio (2018–2023)



3.3 Validating the Granular LLM skills Measure

In this section, we validate our LLM-based (transformer model based) multidimensional measure of traditional IT and AI skills through wage prediction exercises. The primary finding of our analysis is that skill pricing in IT jobs depend on the joint structure of traditional and AI skills and not just on univariate AI intensity measures. To provide empirically robust and theoretically consistent benchmarks for wage-based validation, we construct two measures of AI relatedness used in the related work. We first construct a job-posting-level AI-relatedness measure that replicates the co-occurrence approach used by Babina et al. (2024). Specifically, AI Overlap Measure is defined based on the co-occurrence of skills with four foundational AI pillar domains: Artificial Intelligence, Machine Learning, Computer Vision, and Natural Language Processing. Second, following Acemoglu et al. (2022), we construct two job-posting-level measures of AI intensity. The Narrow AI Index is a binary indicator equal to one if a posting contains any of the 33 AI-related skill keywords from their taxonomy. The Broad AI Index extends this by additionally flagging postings that includes the general terms "AI" or "Machine Learning".

Our construction of the AI Overlap measure follows the identical bottom-up aggregation logic as the original index; it is first computed at the individual skill level and subsequently aggregated.

However, whereas Babina et al. (2024) aggregate these indices to the firm level to capture AI human capital investment, we retain the aggregation at the job-posting level to ensure consistency with the unit of observation used in our zero-shot skill measures and validation framework. We verify that the resulting co-occurrence values closely match those reported in Babina et al. (2024). For example, the AI index score for the skill *TensorFlow* in our data is 0.918, compared to 0.9 in the original study, indicating near-perfect alignment.

We use this job-posting-level AI-overlap measure (Babina et al., 2024) in Table A4 as a univariate benchmark against our 30-dimensional Centered Log Ratio (CLR) skill portfolio. By applying the same train-test split to both the LASSO-based multidimensional model and the linear Babina-style baseline, we directly assess the incremental predictive value of our joint skill representation relative to a univariate AI skill intensity measure.

The first row in Table A4 reports the results for the full sample from 2018 to 2023. We observe that the LLM-based measure reduces mean squared error by 9.5%, and the improvement in out-of-sample R^2 is 74.8% relative to that achieved with the AI overlap measure. Moreover, this analysis is repeated for each year, and we observe an MSE improvement of 7-16%. We also observe that the out-of-sample R^2 with the LLM-based measure is 1.5 to 2 times that achieved with the univariate AI measure.

Next, we do the same comparisons for the Narrow and Broad AI indices using the skill taxonomy in Acemoglu et al. (2022). Table A5 reports the MSE and the out of sample R^2 for these measures. Consistent with the Babina et al. (2024) benchmark, the LLM-based measure demonstrates improvement in out-of-sample predictions than either univariate keyword-based measure.

4 Conceptual Framework

We adopt a nested constant elasticity of substitution (CES) production function to model skill complementarities and substitution for firms, following Krusell et al. (2000); Lindquist (2004). Specifically, we modify the production function in Krusell et al. (2000), where unlike their model we work with one standard capital type and allow three labor skill types (IT Ops, IT Dev, and AI) instead of the two labor types (skilled and unskilled labor). The aggregate production function is given by:

$$Y_t = Z_t \cdot K_t^{\delta} \left[\gamma(O_t)^{\sigma} + (1-\gamma)\left[\alpha(A_t)^{\rho} + (1-\alpha)(D_t)^{\rho}\right]^{\sigma/\rho}\right]^{(1-\delta)/\sigma} \tag{1}$$

where Y_t is aggregate output for the firm at time t, Z_t is total factor productivity (TFP), K_t is physical capital, O_t, D_t, and A_t are IT Operations, IT Development, and AI skill inputs (measured in efficiency units), $\delta \in (0,1)$ is the capital share parameter, $\gamma, \alpha \in (0,1)$ are distribution parameters and $\sigma, \rho \in (-\infty, 1)$, $\sigma, \rho \neq 0$ are substitution parameters between different inputs.

$\Phi_1 = [\alpha(A_t)^{\rho} + (1-\alpha)(D_t)^{\rho}]^{1/\rho}$ is used to define the inner CES aggregate between Development and AI and the outer CES aggregate (Ops skills with Dev-AI aggregate) is defined as $\Phi_2 = [\gamma(O_t)^{\sigma} + (1-\gamma)(\Phi_1)^{\sigma}]^{1/\sigma}$. Using this functional form, the elasticity of substitution between AI and Dev is $\sigma_{A,D} = \frac{1}{1-\rho}$ and Ops and AI is $\sigma_{O,\{D,A\}} = \frac{1}{1-\sigma}$. Similar to Krusell et al. (2000),

when $\rho < \sigma$, the production function exhibits skill complementarity between Development and AI skills relative to Operations skills.[9]

Every skill type is paid its marginal product under competitive labor markets. Consequently, in equilibrium $W_O = \frac{\partial Y}{\partial O}$, $W_D = \frac{\partial Y}{\partial D}$, $W_A = \frac{\partial Y}{\partial A}$ where W_i represents the market wage per efficiency unit of skill type i. The total wage for a worker with skill bundle (O, D, A) is $w = W_O \cdot O + W_D \cdot D + W_A \cdot A$. Further, the CES aggregate Φ_2 exhibits constant returns to scale in the three skill inputs. By Euler's theorem for homogeneous functions we get $O \cdot \frac{\partial Y}{\partial O} + D \cdot \frac{\partial Y}{\partial D} + A \cdot \frac{\partial Y}{\partial A} = (1-\delta) \cdot Y$. For a representative worker, the total wage payment to skills is $w = (1-\delta) \cdot Z_t \cdot K_t^{\delta} \cdot (\Phi_2)^{1-\delta}$.[10] Taking logarithms would give us $\log(w) = \log(1-\delta) + \log(Z_t) + \delta \log(K_t) + (1-\delta)\log(\Phi_2)$.

The nested CES structure yields a nonlinear relationship between log wages and log skill levels $\log(w) = f(\log O, \log D, \log A; \mathbf{\Theta})$ where $\mathbf{\Theta} = \{\delta, \gamma, \alpha, \sigma, \rho\}$ are the CES parameters. To obtain an empirically tractable specification, we employ a third-order Taylor expansion of $f(\cdot)$ around the sample mean log skill levels $(\overline{\log O}, \overline{\log D}, \overline{\log A})$ as follows:

$$
\begin{aligned}
\log(w) \approx & f(\overline{\log O}, \overline{\log D}, \overline{\log A}) \\
& + \sum_{i \in \{O,D,A\}} \frac{\partial f}{\partial \log i}\bigg|_{\bar{\cdot}} [\log i - \overline{\log i}] \\
& + \frac{1}{2} \sum_{i \in \{O,D,A\}} \sum_{j \in \{O,D,A\}} \frac{\partial^2 f}{\partial \log i \, \partial \log j}\bigg|_{\bar{\cdot}} [\log i - \overline{\log i}][\log j - \overline{\log j}] \\
& + \frac{1}{6} \sum_{i \in \{O,D,A\}} \sum_{j \in \{O,D,A\}} \sum_{k \in \{O,D,A\}} \frac{\partial^3 f}{\partial \log i \, \partial \log j \, \partial \log k}\bigg|_{\bar{\cdot}} [\log i - \overline{\log i}][\log j - \overline{\log j}][\log k - \overline{\log k}] \\
& + \text{higher order terms} \qquad (2)
\end{aligned}
$$

where $\big|_{\bar{\cdot}}$ denotes evaluation at sample means.

The nested CES structure, combined with the taylor series approximation, provides testable predictions about skill complementarities. First, if $\rho < 1$ (low elasticity of substitution in the inner nest), Development and AI skills are complements, implying $\gamma_{DA} > 0$ when Operations skills are at their sample mean. Second, if $\rho < \sigma$ (lower elasticity in the inner nest relative to the outer nest), AI skills are more complementary with Dev than with Ops, implying $\gamma_{OA} < 0$ when Development skills are at their sample mean. Third, if skill synergies are amplified when all three skills are present together, we predict $\delta_{ODA} > 0$, indicating that the complementarity between any two skills is stronger when the third skill is also present at higher levels.

[9] In Krusell et al. (2000) the equivalent of AI is equipment capital, for Development it is skilled labor, and for Operations it is unskilled labor.

[10] We drop the subscript t here to keep the notation simple.

5 Identification and Empirical Strategy

In this section we introduce notation for our empirical specification motivated from the previous section. Each job posting is denoted by $i \in \{1, 2, \ldots, n\}$, in firm $j \in \{1, 2, \ldots, J\}$, state $s \in \{1, 2, \ldots, S\}$, in month-year $t \in \{1, 2, \ldots, T\}$. Precisely, for each job posting i, we construct skill scores as $s_{ik} = \sum_{j \in \mathcal{C}_k} p_{ij}$ where $k \in \{\text{O, D, A}\}$ and $\mathcal{C}_k$ denotes the set of skills in category k (see Section 3 for the complete list and aggregation details).

We cannot use the raw skill scores directly as covariates, since they are compositional data that only contain relative information (Aitchison, 1982; Greenacre, 2021). An increase in one skill category mechanically implies changes in the relative importance of other categories. Consequently, regression coefficients cannot be interpreted in the standard way, as an increase in one skill does not hold other skills constant in relative terms. To obtain a representation of skill portfolio in real space and ensure that coefficients capture shifts in relative skill intensity, we apply the centered log-ratio (CLR) transformation.

The CLR transformation treats each category symmetrically by centering on the geometric mean. This is preferable to additive log ratios (ALRs), which require choosing an arbitrary baseline category which in turn will impact the interpretation of the estimated coefficients in the regression model. Moreover, CLR is easier to calculate and interpret than isometric log ratios (ILRs), which involve a complex partitioning step. For each category k, we define $\text{CLR}_{ik} = \log(s_{ik}) - \frac{1}{3}\sum_{k=1}^{3} \log(s_{ik}) = \log\left(\frac{s_{ik}}{g_i}\right)$ where $g_i = \left(\prod_{k=1}^{3} s_{ik}\right)^{1/3}$ is the geometric mean of the three skill scores. By construction, CLR coordinates sum to zero. This creates perfect collinearity, so the regression specification removes one of the categories, which henceforth we will call the baseline category (CLR:AI). Our primary specification with main effects, two-way and three-way interactions is:

$$\begin{aligned}\log(w_{ijst}) = &\sum_{k=1}^{2} \beta_k \text{CLR}_{ijk} + \sum_{k=1}^{2}\sum_{\ell>k}^{3} \gamma_{k\ell}(\text{CLR}_{ijk} \times \text{CLR}_{ij\ell}) \\ &+ \delta_{ODA}(\text{CLR}_{ijO} \times \text{CLR}_{ijD} \times \text{CLR}_{ijA}) + \mathbf{X}'_{ijst}\boldsymbol{\theta} + \alpha_j + \mu_s + \lambda_t + \epsilon_{ijst}\end{aligned} \tag{3}$$

where w_{ijst} is the wage for job posting i at firm j in state s at time t. Fixed effects include α_j (firm), μ_s (state), and λ_t (year-month). The vector $\mathbf{X}_{ijst}$ includes education requirements, experience requirements, employment type, and remote work indicators. Standard errors are clustered at the firm level. This maps to the Taylor expansion of the wage model in section 4. Note that all interaction parameters are separately identified as they introduce higher-order polynomial terms not collinear with main effects.

5.1 Identification

Our identification assumption is that, conditional on firm, state, and time fixed effects (month-year) and posting-level covariates, variation in skill requirements across postings is uncorrelated with the error term $\mathbb{E}[\epsilon_{ijst}|\text{CLR}_{iO}, \text{CLR}_{iD}, \text{CLR}_{iA}, \mathbf{X}_{ijt}, \alpha_j, \mu_s, \lambda_t] = 0$. The parameters of inter-

est are identified from within-firm, within-state, within-time variation in skill composition across job postings. To see the source of the identification, we decompose the CLR-transformed skill measure into a conditional mean and residual $\text{CLR}_{ijk(st)} = \mathbb{E}[\text{CLR}_{ik}|\alpha_j, \mu_s, \lambda_t] + \widetilde{\text{CLR}}_{ijk(st)}$, where $\widetilde{\text{CLR}}_{ijk(st)} = \text{CLR}_{ijk(st)} - \mathbb{E}[\text{CLR}_{ik}|\alpha_j, \mu_s, \lambda_t]$ is the deviation from the conditional mean. Identification for the parameters requires sufficient variation in skill composition within the fixed effect groups.

We note that our use of firm- and year-month-fixed effects is different from recent critiques of two-way fixed effects in staggered difference-in-differences designs (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021). These critiques apply to settings in which binary treatments are assigned at different times across units, leading to negative weights and biased estimates of the treatment effect. We do not estimate treatment effects from policy changes.

5.2 Interpretation of Parameters

In the model with two-way and three-way interactions, the marginal effect of skill category k on wages is $\frac{\partial \log(w_{ijst})}{\partial \text{CLR}_{ijk}} = \beta_k + \sum_{\ell \neq k} \gamma_{k\ell} \cdot \text{CLR}_{ij\ell} + \delta_{ODA} \cdot \text{CLR}_{ijm} \cdot \text{CLR}_{ijn}$ where $m, n \neq k$ denote the two other skill categories. The coefficient β_k represents the marginal wage return to skill category k when all other CLR-transformed skills equal zero (i.e., at their geometric mean). Since this configuration may not represent typical job postings in our sample, we report marginal effects evaluated at meaningful skill compositions (e.g., sample mean, specific percentiles) in our results section.

In a fully interacted model, the two-way interaction coefficients $\gamma_{k\ell}$ cannot be interpreted in isolation. The complementarity or substitutability between skills k and ℓ depends on the level of the third skill. Specifically, the cross-partial derivative measuring the interaction between skills k and ℓ is $\frac{\partial^2 \log(w_{ijst})}{\partial \text{CLR}_{ijk} \partial \text{CLR}_{ij\ell}} = \gamma_{k\ell} + \delta_{ODA} \cdot \text{CLR}_{ijm}$ where m denotes the third skill category not indexed by k or ℓ. Thus, $\gamma_{k\ell}$ represents the complementarity between skills k and ℓ only when the third skill is at its geometric mean ($\text{CLR}_{ijm} = \log(\frac{g_{ij}}{g_{ij}}) = 0$). For example, γ_{DA} captures AI-Development complementarity when Operations skills are at their geometric mean. Further, a positive three-way interaction ($\delta_{ODA} > 0$) indicates that skill synergies are amplified when all three skills are present.

6 Results

6.1 Primary Specification

The outcome variable wages (posted salaries) is expressed in logs, and we transform the three aggregated skills (Operations, Development, and AI) from LLM-generated probability scores into CLRs. Table 1 provides the results with the main effects only without interactions in column (1), and then we gradually add two way interactions (Column (2)) and finally the three way interactions in Column (3). Each specification includes firm-level, state-level, and year-month level fixed effects, and the standard errors are clustered at the firm level. We also include a set of controls. We include dummy variables for the educational degree required in the job ad, minimum

and maximum experience, work arrangement (remote, in-person, or hybrid), an internship dummy, and employment type (full-time and part-time).

Table 1: Wage Premia using CLR Transformations

	Dependent variable: Log(Salary)		
	Main Effects	2-Way Interactions	3-Way Interactions
	(1)	(2)	(3)
CLR: Operations	0.046	0.157***	0.137***
	(0.061)	(0.045)	(0.046)
CLR: Development	0.318***	0.543***	0.536***
	(0.097)	(0.090)	(0.085)
Operations × Development		0.149***	0.206***
		(0.017)	(0.033)
Operations × AI		−0.296***	−0.233***
		(0.034)	(0.032)
Development × AI		0.231***	0.342***
		(0.055)	(0.060)
Operations × Development × AI			0.189***
			(0.054)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	870,797	870,797	870,797
R^2	0.531	0.538	0.538
Adjusted R^2	0.530	0.537	0.537

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Controls include education level, years of experience, employment type, remote-work status, and an indicator for internship. Standard errors clustered at the firm level are reported in parentheses.

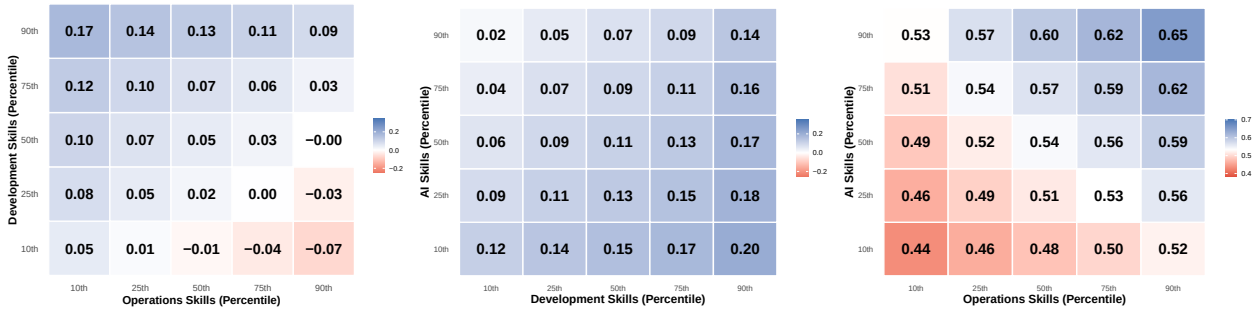
In Column (1) in Table 1 we observe that relative to the baseline of AI CLR there is a positive premium for both IT Operations and Development. Moreover, the magnitude for IT Development is much higher and the 95% confidence interval does not include 0. However, the premium to IT Operations is not significant at the conventional level of significance (5%). Column (2) presents the results with two way interactions of AI with IT Operations and Development. We find that the main effects on Operations and Development relative to the baseline for AI are positive, and that AI skills are complementary to IT development. However, it suggests a substitution between Operations and AI. Finally, we present the results of the fully interacted 3 way interactions in Column (3). We find that main effects and two way interactions are in the same direction as Columns (1) and (2). However, there appears to be a premium to having all three skills.

We calculate the marginal effect of AI on wages as $\frac{\partial \log(w)}{\partial \text{CLR AI}} = -0.233 \cdot \text{CLR Ops} + 0.342 \cdot \text{CLR Dev} + 0.189 \cdot \text{CLR Ops} \cdot \text{CLR Dev}$ where the coefficients correspond to the two-way interactions (Operations × AI and Development × AI) and the three-way interaction (Operations × Development × AI) from Column (3) in Table 1. The marginal effect of AI is a function of the other two CLRs. Therefore, we vary the Operations and Development CLR between their values

of 10th, 25th, 50th, 75th and 90th percentile and calculate the marginal effect of AI for each pair of these combinations (25 in total).

Figure 3 Panel (a) provides a heatmap to visualize these marginal effects. The heatmap reveals three key patterns.[11] First, the heat map suggests substitution between Operations and AI. Moving from the 10th to the 90th percentile of CLR Operations reduces the marginal effect, with the marginal effect declining from +0.05 to −0.07 in the bottom row (low Development). This translates from a 5.13% wage premium to a wage penalty of 6.76% where the % change from log points is computed as $(e^{\beta} - 1) \times 100$. Second, the heatmap suggests complementarity between AI and Development. A move from the 10th to the 90th percentile of Development increases the marginal effect of AI, with values rising from 0.05 to 0.17 in the left most column (low Operations). This suggests the marginal returns to AI increases from 5.13% to 18.53% for a skill portfolio that combines high levels of Development holding Operations at the 10th percentile. Third, and most importantly, the positive three-way interaction (+0.189) manifests as non-parallel rate changes across the grid. Specifically, when moving from low to high Operations, the decline in marginal effect is smaller at high Development levels (−0.08, from 0.17 to 0.09 in the top row) than at low Development levels (−0.12, from 0.05 to −0.07 in the bottom row). This 0.04 difference in log points (or 4.08%) demonstrates that the substitution effect of Operations and AI is attenuated when Development skills are also present. Moreover, the premium of Operations with Development on the marginal effect of AI is evident from the diagonal when both skills are moved from the 10th to the 90th percentile. This diagonal move suggests that we move from 5.13% wage premia to 9.42% when both Development and Operations are moved together from the 10th to the 90th percentile.

Figure 3: Marginal Effect of AI, Operations and Development on Posted Salaries



a. AI b. Operations c. Development

Next, we present the heatmaps for the marginal effects of IT Operations and IT Development. Panels (b) and (c) in Figure 3 provide these graphs. Panel (b) in Figure 3 displays the marginal effect of Operations skills on log salary across different combinations of CLR Development and CLR AI skill levels. The marginal effect is calculated as $\frac{\partial \log(w)}{\partial \text{CLR Ops}} = 0.137 + 0.206 \cdot \text{CLR Dev} - 0.233 \cdot$

[11] All three skills are expressed in CLRs. Consequently, increasing skill k corresponds to an increase in skill k relative to the geometric mean.

CLR AI + 0.189 · CLR Dev · CLR AI, with values ranging from 0.02 to 0.20 across the skill space. Operations skills exhibit universally positive returns, though the magnitude varies systematically with complementary skills. The vertical gradient reveals the AI substitution effect. When AI skills are low (10th percentile, bottom row), Operations commands the highest marginal returns of 0.12 to 0.20 in log points (or 12.75% to 22.14%), but as AI increases to the 90th percentile (top row), these returns diminish substantially to 0.02 to 0.14 (or 2.02% to 15.03%). This pattern confirms that AI and Operations are substitutes, as AI tools likely automate routine IT operations, thereby decreasing the incremental value of Operations expertise. The horizontal gradient demonstrates Development complementarity. Operations returns are lowest at 0.02 to 0.12 when Development is at the 10th percentile (left column) and increase to 0.14 to 0.20 when Development reaches the 90th percentile (right column), representing an approximately 0.08 log-point increase. Moreover, the wage penalty from high AI on Operations returns is attenuated when Development skills are elevated as well (positive three-way interaction).

Panel (c) in Figure 3 displays the marginal effect of Development skills on log salary across different combinations of Operations and AI skill levels. The marginal effect is calculated as $\frac{\partial \log(w)}{\partial \text{CLR Dev}} = 0.536 + 0.206 \cdot \text{CLR Ops} + 0.342 \cdot \text{CLR AI} + 0.189 \cdot \text{CLR Ops} \cdot \text{CLR AI}$, with values ranging from 0.44 to 0.65. Development skills command uniformly positive and substantial returns across all skill combinations, demonstrating that software development capabilities are universally valued in the IT labor market regardless of other skills possessed. Development returns increase from 0.44 to 0.53 when AI is at the 10th percentile (bottom row) to 0.52 to 0.65 when AI reaches the 90th percentile (top row), representing wage premia of approximately 0.09 to 0.13 log points (9.42% to 13.88%). This suggests that AI and Development are complements. The horizontal gradient shows the Operations complementarity. Development returns rise by approximately 0.08 log points as Operations increases from the 10th to 90th percentile across all rows, reflecting the positive coefficient on the IT Operations × IT Development interaction term. The three-way interaction creates a multiplicative synergy most visible in the upper-right corner where the marginal effect reaches 0.65.

Finally, Figure A3 visualizes the estimated wage contribution from skill portfolios across all three dimensions of IT human capital where each point in the three-dimensional space represents a unique skill portfolio combination, with position determined by Ops (x-axis), Dev (y-axis), and AI (z-axis) skill levels, and color intensity indicating the estimated log wage contribution from that skill bundle. The visualization reveals several key patterns in the wage structure. The highest estimated wages (blue region) concentrate where all three skills are elevated, particularly when Development and AI skills are both high. Examining the vertical gradients along the AI dimension (z-axis), the rate at which wages increase with AI skills varies systematically with Development skill. When Development is high (positive y-values), increasing AI generates steeper wage gains, whereas when Development skills are low (negative y-values), AI additions yield smaller returns, or even wage decreases, due to the substitution effect with Operations skills.

6.2 Robustness Checks

Next, we conduct a series of robustness checks to alleviate concerns about identification and measures.

Labor market dynamics. The primary specification in Table 1 uses fixed effects at the level of firm, month-year, and state fixed effects. State fixed effects are included to account for time-invariant regional unobserved heterogeneity that may affect wage premia. However, state fixed effects may be too coarse to capture region-specific wage differences that are likely correlated with the skill portfolio. Instead of state fixed effects, we estimate the primary specification using MSA fixed effects. This would account for concerns that different MSAs have varying unobserved labor market dynamics. For instance, labor markets in California, such as San Francisco (SF), likely differ from those in Sacramento due to a higher concentration of AI-intensive firms in SF. If tech hubs systematically post higher wages due to factors such as geographic sorting of high-growth firms, which may or may not be related to skill portfolios, our estimated parameters could conflate the skill premium with region-specific wage premia. Table A6 in the Appendix presents the results with MSA fixed effects. We observe very similar estimated parameter values for the main and interaction effects, which alleviates concerns about geographic confounding.

Further, to ensure that our results are not confounded by varying levels of AI exposure across different MSAs, we construct MSA level leave one out measure of AI exposure. If MSAs with higher AI exposure have a higher concentration of firms with larger wage premia that are not explained solely by the skill portfolio but by unobserved factors, such as differential living conditions or higher wage premia due to greater competition for talent, then this will likely confound the skill portfolio premia. We construct a new measure of AI exposure at the MSA level using Babina et al. (2024)'s AI overlap index. This measure varies at the MSA month-year level, hence it will be identified in a specification with MSA fixed effects. However, we construct it as a leave-one-out mean to remove any mechanical correlation with the ith firm.

For each firm j in MSA m at time t, we construct: $\text{AI}_{-j,mt} = \frac{1}{N_{mt}-1} \sum_{k \neq j, k \in m} \text{AI}_{kt}$ where N_{mt} is the number of firms in MSA m at time t and AI_{kt} is the AI overlap score. Table A7 reports the results for this analysis.[12] Column (1) reports the parameter estimates with firm, month-year, and MSA fixed effects for baseline comparison, and then Column (2) adds the MSA AI exposure. We see that the results are comparable across the two specifications for both main effects and interactions, in terms of the magnitudes of the point estimates and their standard errors.

Quantile Regression - Effects Across the Wage Distribution. To assess if the complementarities and substitution effects of skill portfolios that we observe in the fully interacted specification are concentrated across the entire wage distribution or in certain quantiles, we estimate quantile regressions at the 25th, 50th, and 75th quantiles. Because conditional quantiles are not linear

[12] We note that the sample in this specification differs slightly from our primary sample for two reasons. First, we remove observations with missing values in the AI exposure index. Second, we cannot construct this measure for cases in which an MSA month cell contains only one firm or for job postings where the MSA information is missing.

operators, estimating panel regressions with high-dimensional fixed effects cannot be done in the same way as in linear models. For this, we use the results in Canay (2011), which establishes that if the fixed effects are "location shifters," they can be eliminated using a two-step procedure.

Following Canay (2011), we model wages as $\log(w_{ijst}) = \mathbf{CLR}'_{ist}\boldsymbol{\theta}(U_{ijst}) + \alpha_j + \mu_s + \lambda_t + \epsilon_{ijst}$ where $U_{ijst} \sim \text{Uniform}[0,1]$ represents the job posting's position in the residual wage distribution, $\boldsymbol{\theta}(\tau)$ captures skill effects at quantile τ, and $(\alpha_j, \mu_s, \lambda_t)$ are firm, state, and year-month fixed effects. Canay (2011)'s key assumption is that fixed effects are *location shifters*, meaning that they affect all quantiles uniformly. This implies $Q_\tau(\log(w_{ijst}) \mid \mathbf{CLR}, \mathbf{X}, \alpha_j, \mu_s, \lambda_t) = \mathbf{CLR}'_{ist}\boldsymbol{\theta}(\tau) + \alpha_j + \mu_s + \lambda_t$ where $Q_\tau(Y \mid \cdot)$ denotes the τ-th conditional quantile function. Under this assumption, the fixed effects can be consistently estimated from the conditional mean $\mathbb{E}[\log(w_{ijst}) \mid \mathbf{CLR}, \mathbf{X}, \alpha_j, \mu_s, \lambda_t] = \mathbf{CLR}'_{ist}\boldsymbol{\theta}_\mu + \alpha_j + \mu_s + \lambda_t$ where $\boldsymbol{\theta}_\mu = \mathbb{E}[\boldsymbol{\theta}(U_{ijst})]$. We implement this via Frisch-Waugh where for each variable Z, we compute residuals $\tilde{Z}_{ijst}$ from regressing Z on firm, state, and year-month fixed effects plus controls. Quantile regression on $\{\tilde{Z}\}$ then consistently estimates $\boldsymbol{\theta}(\tau)$. Table A8 reports the results of the quantile regression. We see similar patterns of complementarity (Development and AI) and substitution (Operations and AI) across all three quantiles.[13]

Firm Age. Our results of skill portfolio on wages can be confounded if younger and older firms have unobserved systematic differences that are correlated with both skills and posted salaries. Younger publicly traded firms (e.g., recent IPOs) may systematically post jobs requiring AI and development skills together, while also willing to pay high wage premia for factors unrelated to skill productivity. These can include aggressive growth strategies, risk levels, competition for talent from established tech firms, or differences in the compensation structure after the IPO. To alleviate these concerns, we create a firm age variable as the difference between the founding year and the job posting year. The founding year for each firm is obtained from the Capital IQ database. Next, we create separate samples based on this firm age variable and estimate our three-way interaction model. Table A9 reports these results. Column (1) provides the baseline estimates for the full sample. Note for this analysis, we do lose a few observations due to missing or outlier values in the founding year variable. Columns (2) and (3) provide the results separately for younger and older firms, respectively, where we split at the median. We see comparable estimates of complementarity and substitutability between the three skill portfolios across these two columns, alleviating concerns that firm maturity is confounding the skill premia findings.

Alternative Definitions of Skill Portfolio. We test the robustness of our results to alternative skill portfolio definitions. We create two scenarios. In the first scenario, we move database management to IT Development, and in the second, we move cloud computing and dev-ops into IT Operations. We estimate the fully interacted model (a three-way interacted specification) with these alternative definitions. Table A10 provides the baseline in column (1) (same as Column (3) Table 1), and columns (2) and (3) provide the estimates with the alternative definitions. The pa-

[13]The quantile regression on residualized variables is implemented using `quantreg` package in R.

rameter estimates for the main and the interaction effects are robust to these alternative definitions of IT Operations and IT Development.

Sample Restrictions. Finally, we test the robustness of our results to alternative sample restrictions. The first check is to remove the top 1% and bottom 1% of job postings in the salary distribution to make sure the results are not driven by outliers. Table A11 reports these results. Column (1) reports the baseline with the full sample, and column (2) reports the results removing the top and bottom percentiles. We see that the estimated parameters are comparable in magnitude and standard errors across the two columns. Next, we consider another type of sample restriction. We remove firms with very few postings. Our primary specification includes firm-fixed effects, and we need substantial within-firm variation to estimate the parameters of interest precisely. Columns (3) and (4) in Table A11 report the estimates restricting to firms with at least 5 postings and 10 postings, respectively. This restriction does change the sample size, but the estimates remain robust to our baseline model with the full sample.

7 Heterogeneous Impact of Skill Type on Wages

In this section, we examine heterogeneity in the wage premia for IT skill portfolio by experience, work arrangement type, and job location within tech hubs.

Experience We begin by splitting the sample by the required experience level in the job posting. We construct three groups. Junior hire with 0-2 years of experience, mid-tier hire with 3-7 years of experience and finally senior hire with atleast 8 years of experience.[14] The wage premium to the three types of IT skills likely vary by experience as the senior roles often require non technical skills such as managerial and cognitive skills. Consequently, we hypothesize that the complementarity premium between the skills is highest for the junior roles and the premium declines with seniority.

Table 2 reports the results for the three groups separately in Columns (1)–(3). As for the coefficients on the interactions, there appears to be a positive premium for all three skills only among entry-level and mid-tier workers, and none among senior workers. The premium of combining Development skills with AI appears to be driven primarily by entry- and mid-tier jobs, with the highest premium for entry-level jobs.

However, in a fully interacted model, one cannot interpret the coefficients on higher-order interactions in isolation. We use the same calculation as before, in which we estimate the marginal effect of AI on wages as a function of both Operations and Development skills, since they enter the marginal effect formula through higher-order interactions. We consider a grid of values for Operations and Development. We extract the percentiles for these two skills and calculate the

[14] While constructing the covariates, about 29% of postings did not explicitly mention minimum years of experience required. We include a missing experience dummy for the main regressions. However, for the heterogeneity analysis, the sample without any minimum experience required is added to 0-2 years to keep the samples comparable to the main analysis.

Table 2: Wage Premia by Experience Level: Three-Way Skill Interactions

	Dependent variable: Log(Salary)		
	Entry (0-2 yrs)	Mid (3-7 yrs)	Senior (8+ yrs)
	(1)	(2)	(3)
CLR: Operations	0.195***	0.035	0.057*
	(0.063)	(0.030)	(0.029)
CLR: Development	0.659***	0.323***	0.184***
	(0.107)	(0.041)	(0.036)
Operations × Development	0.201***	0.184***	0.085**
	(0.035)	(0.023)	(0.034)
Operations × AI	−0.281***	−0.195***	−0.155***
	(0.043)	(0.019)	(0.052)
Development × AI	0.384***	0.168***	0.003
	(0.066)	(0.030)	(0.106)
Operations × Development × AI	0.156**	0.119***	−0.030
	(0.072)	(0.045)	(0.101)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	367,750	369,784	133,263
R^2	0.551	0.452	0.451
Adjusted R^2	0.548	0.448	0.445

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Controls include education level, years of experience, employment type, remote work status, and an indicator for internship. Note that there are 29% job postings do not explicitly require a minimum number of years of experience. This is part of column (1) sample entry-level experience for the heterogeneity analysis. Standard errors clustered at the firm level are reported in parentheses.

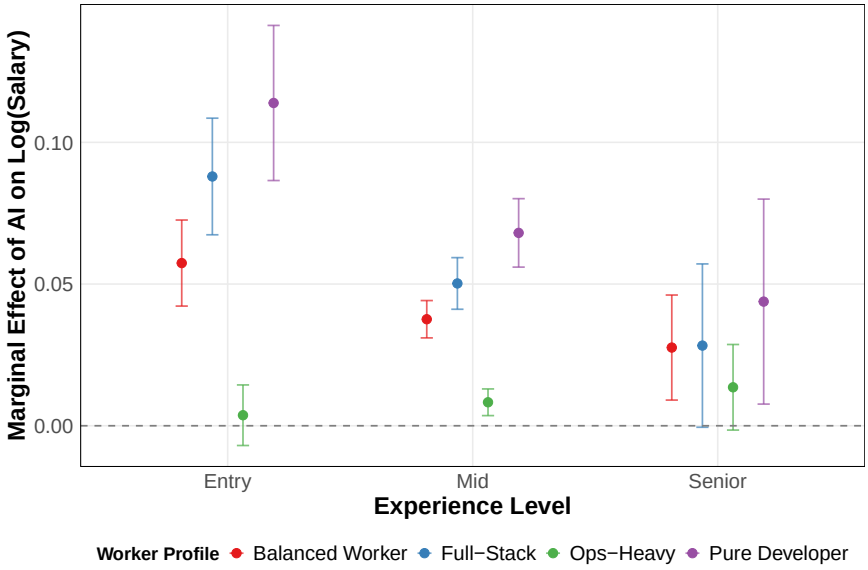
marginal effect of AI on wages for every pair of these two skill percentiles. To implement this analysis, we consider four portfolios. We call these pure developer (High Dev (p75), Low Ops (p25)), balanced worker (Medium Dev (p50), Medium Ops (p50)), Ops heavy (Low Dev (p25), High Ops (p75)), and full stack (High Dev (p75), Medium Ops (p50)). Figure 4 shows the marginal effects for AI skills by experience and worker skill portfolios. We gain three key insights. First, the marginal effect of AI skills is highest for entry-level jobs across all portfolios, except for Ops-heavy roles, relative to mid-tier and senior jobs. Second, this premium increases with higher Development skills, peaking for pure developers (entry- and mid-tier). Third, this skill complementarity is absent among job postings requiring senior workers.

In order to gain deeper insights into this result, we construct posting-level measures of non-technical skill demand following Deming and Kahn (2018). Specifically, we map Lightcast skill tags to four non-technical categories: People Management, Project Management, Cognitive Skills, and Social Skills. For each posting, we compute a raw count of skill tags within each category, independently of the transformer-based IT skill measures used in the primary analysis. Table A12 in the Appendix reports the mean raw counts and standard errors for these four categories across three experience levels.

Table A12 reveals a clear and monotonic shift in non-technical skill demand with seniority, providing a direct mechanism for the pattern in Table 2. People Management requirements nearly

triple from entry to senior roles (0.365 to 0.982), Social skill demands rise by 60% (0.697 to 1.115), and Cognitive skill demands increase by 45% (0.910 to 1.323), with Project Management doubling over the same career progression (0.034 to 0.071). These shifts indicate that senior IT roles are fundamentally different in their human capital composition: as workers advance, the wage-relevant skill bundle migrates away from technical execution, where AI-Development complementarities are strongest, toward coordination, leadership, and judgment-intensive tasks that are less susceptible to AI augmentation (Brynjolfsson et al., 2025; Deming, 2017). The marginal contribution of the technical skill portfolio to wage determination therefore weakens at senior levels not because AI becomes less productive, but because wages increasingly reflect non-technical capabilities that are orthogonal to the AI-Development interaction. This finding is consistent with Deming (2017), who shows that tasks requiring interpersonal flexibility and tacit judgment remain the most durable barriers to automation, and with Brynjolfsson et al. (2025), who documents that AI primarily reshapes early task bundles within careers rather than senior coordination roles.

Figure 4: Marginal Effect of AI Skill on Wages by Experience



Work Arrangement Next, we explore heterogeneity by work arrangement. Note that we consider data from 2018 to 2023, including the onset of remote work during the 2020 pandemic. In our sample, job postings mention three types of worker arrangements. Such remote arrangements are much more likely in IT jobs than in other occupations. This includes the largest proportion of on-site/in-person jobs (73.1%), 23.8% remote jobs, and the remaining jobs that allow hybrid in-person and remote work arrangements.[15] Table 3 presents the results of the three-way interaction model, disaggregated by work arrangement. The substitution and the complementarity patterns show up

[15]The "on-site" sample for the heterogeneity analysis is constructed using the remote: none and remote: stated none categories, where the former captures job postings without any mention of remote work, and the latter are postings that explicitly mention remote work is not allowed.

Table 3: Wage Premia by Work Arrangement: Three-Way Skill Interactions

	Dependent variable:Log(Salary)		
	On-site	Remote	Hybrid
	(1)	(2)	(3)
CLR: Operations	0.176***	−0.030*	−0.032
	(0.045)	(0.018)	(0.069)
CLR: Development	0.589***	0.283***	0.273***
	(0.080)	(0.040)	(0.070)
Operations × Development	0.198***	0.219***	0.261***
	(0.032)	(0.026)	(0.049)
Operations × AI	−0.259***	−0.133***	−0.125*
	(0.035)	(0.018)	(0.064)
Development × AI	0.333***	0.284***	0.261***
	(0.068)	(0.045)	(0.080)
Operations × Development × AI	0.144**	0.283***	0.365***
	(0.057)	(0.041)	(0.102)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	636,928	211,622	22,247
R^2	0.533	0.604	0.619
Adjusted R^2	0.531	0.601	0.603

Notes: *p<0.1; **p<0.05; ***p<0.01. Reference category: CLR AI (omitted). Remote stated none has been combined with remote none for constructing the "on-site" sample. Controls include education level, years of experience, employment type, and an indicator for internship. Standard errors clustered at the firm level are reported in parentheses.

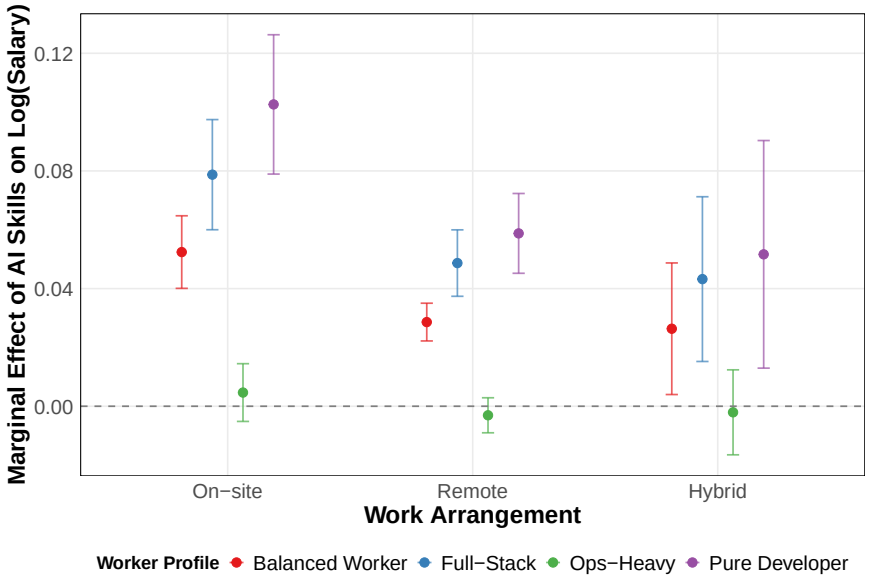
in all three groups although with varying point estimates.

As in the analysis above, we calculate the marginal effect of AI on log wages across four worker portfolios as used in the previous section on experience. Figure 5 shows the marginal effects separately for the skill portfolios and worker arrangement. This indicates a positive marginal effect of AI on wages across all work arrangements and worker types, except for Ops-heavy workers. However, these AI premia are highest when combined with strong Development skills. Because the sample size for hybrid workers is small, we observe larger standard errors for this group than for the other two groups.

Tech Hubs. Next, we examine the heterogeneity of skill demand for wage premia by whether the job location is in a tech hub. This is critical because labor market characteristics likely vary substantially across tech hubs relative to other regions. For this analysis, we obtained MSAs for the job location from the Lightcast dataset. We identify five major tech hubs: San Jose, SF, Seattle, Austin, and Boston. Additionally, we classify four MSAs as secondary tech hubs: Denver, Atlanta, NYC, and Washington, DC. Together, these nine MSAs will be called "tech hubs," and the remaining will be called "non-tech hubs".

Table A13 shows the regression results separately done for the job postings in tech hubs and non-tech hubs in Columns (1)-(2). We observe that the overall direction of the coefficients remains consistent with the results from our full-sample specification. However, there appear to be differ-

Figure 5: Marginal Effect of AI Skill on Wages by Worker Arrangement

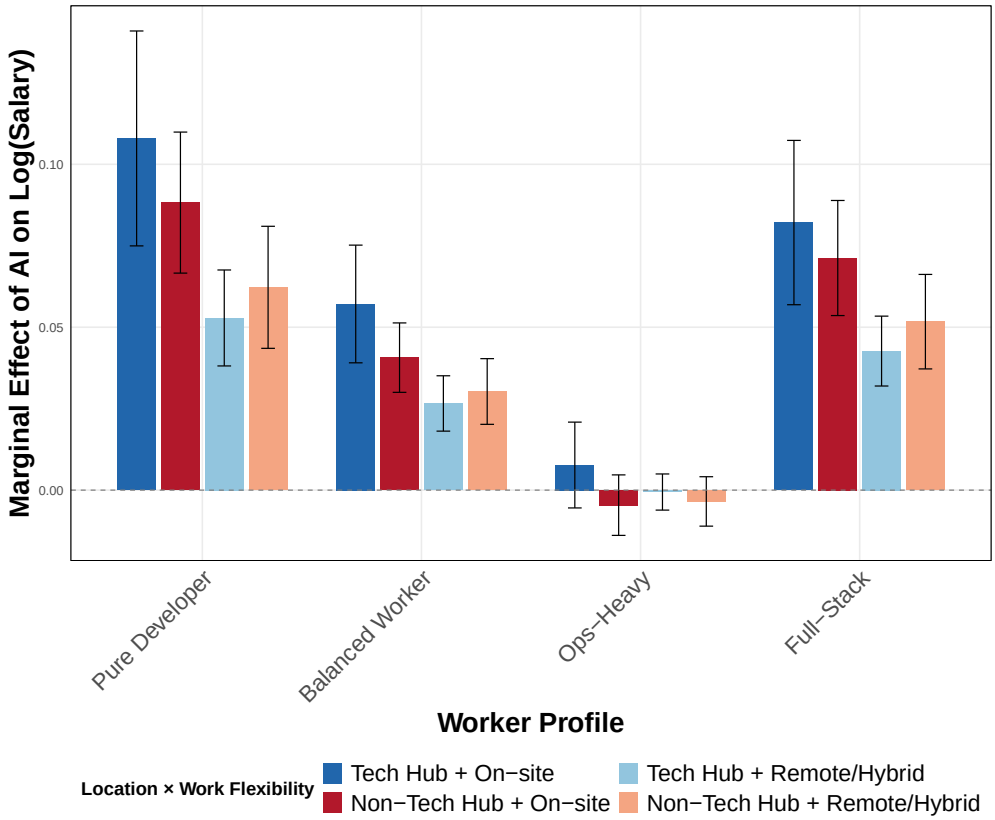


ences in the magnitudes of the higher-order interactions between tech and non-tech hub jobs. To shed more light on the AI premium for wages, we create two worker profiles: i) Dev heavy profile (75th percentile), low Operations (25th percentile), and ii) Ops heavy profile (25th percentile), low Development (75th percentile). Using these, we calculate the marginal effect of AI and display it using a gradient scale in Panels (a) and (b) in Figure A4. Panel (a) shows that, for development-focused workers, the AI wage premium is higher for jobs listed in tech hubs than in non-tech hubs. On the contrary, the AI premium for the Ops-heavy profile is negligible in both types of areas, as evidenced by very small marginal effects in panel (b) in terms of magnitude.

Work Arrangement and Tech Hubs. In this section, we examine whether the wage premia of skill portfolios are affected by both work arrangements and location. We test using a 2×2 design whether premia differ by work flexibility for jobs in tech and non-tech hubs. The work arrangement is grouped into two: the first group comprises jobs that allow remote or hybrid arrangements, and the other group comprises postings that either explicitly state that the job is not remote or do not mention remote arrangements. The other grouping is whether the job is located in a tech hub, uses the same definition of tech hubs as in the previous analysis. The estimated parameters are reported in Table A14 in the Appendix. The substitution and complementarity patterns appear in every column, as in our main results.

Figure 6 shows that the marginal effect of AI on wages is positive and statistically significant at the conventional levels only for pure developers, balanced workers, and full-stack developers. However, the marginal effect is highest for on-site workers in tech hubs relative to those in non-tech hubs. However, the marginal effects for remote/hybrid work are higher in non-tech hubs, although the differences with remote/hybrid in tech hubs are small.

Figure 6: AI Wage Premium Across Location and Work Flexibility



8 Discussion and Conclusion

This study uses large language models to construct multivariate, data-driven measures of AI skill demand and to estimate wage premia arising from complementarities between emerging AI skills and established IT expertise. The study highlights the limitations of univariate AI exposure measures and static skill taxonomies commonly used in research on technology and labor (Acemoglu et al., 2022; Alekseeva et al., 2021; Babina et al., 2024; Deming and Kahn, 2018; Deming and Noray, 2020; Mäkelä and Stephany, 2024). By extracting latent, multidimensional skill demand and embedding these measures within a compositional framework, we offer a scalable approach to studying technological change in settings characterized by rapid innovation and implicit skill requirements.

Managerial Implications This study advances the task-based literature on skill-biased technological change (SBTC) by demonstrating that AI's impact is governed by complex synergies rather than simple occupation-level displacement. We provide empirical evidence that the market price of IT human capital in the AI era is a function of the joint structure of traditional and emerging capabilities. This shifts the academic narrative from univariate "AI exposure" to "portfolio complementarity," which helps explain why aggregate employment effects may appear muted despite

substantial restructuring at the task level (Acemoglu et al., 2022; Bessen, 2019).

Our results show that while AI exhibits strong complementarities with IT Development skills, its interaction with IT Operations skills is negative, so that "pure" operations roles are at the highest risk of stagnation or displacement, as AI tools increasingly automate routine system monitoring and troubleshooting (Acemoglu and Restrepo, 2018; Autor and Dorn, 2013). For managers, this implies that hiring strategies should prioritize "Full-Stack" or "Dev-AI" profiles over standalone systems administrators to avoid over-investing in routine roles that are rapidly being automated. Moreover, it calls for managers to consider reallocating and retraining IT Operations staff in line with an option we discuss below.

Importantly, this negative interaction does not imply a uniform devaluation of Operations expertise. Rather, AI reshapes the relative contribution of Operational skills, reducing their marginal wage association in AI-intensive environments. Consistent with this interpretation, we observe a positive three-way interaction among AI skills, IT Development skills, and IT Operations skills, suggesting that operations expertise adds economic value when embedded within broader, integrative skill bundles. Workers who can design systems while accounting for operational constraints benefit from AI, even as routine execution tasks are automated. From an organizational design perspective, firms should restructure IT departments to align with integrated product teams that combine architectural design with operational oversight, rather than maintaining siloed infrastructure units. These findings are consistent with recent industry reports, such as PwC AI Jobs Barometer Report (2025), and with Anthropic AI's internal research (Huang et al., 2025), suggesting that engineers can focus on high-level design, planning, and oversight while AI can handle routine coding and debugging tasks.

A critical finding from our analysis is the heterogeneity of AI returns across career stages. We further find that the demand for non-technical skills, such as people management, project management, social and cognitive skills, is higher in senior roles than in junior and mid-career roles. The concentration of AI development wage premia among junior and mid-career workers suggests that AI primarily reshapes early task bundles within jobs, reallocating effort from execution toward design, integration, and system-level problem-solving, with muted marginal effects for senior, coordination-intensive roles (Brynjolfsson et al., 2025). For policymakers and HR leaders, this suggests that public reskilling funds and internal training budgets should be allocated disproportionately toward junior-to-mid-career technical staff, as they represent the cohort with the highest marginal wage and productivity sensitivity to AI integration.

The spatial heterogeneity in our results indicates that AI-driven task reallocation is mediated by local labor market thickness and knowledge spillovers, with dense technology regions better able to recombine AI with existing IT capabilities into high-value Development focused task bundles. This pattern indicates that the returns to AI skills extend beyond individual workers' capability and depend on regional ecosystems that facilitate recombination, learning, and the deployment of complementary technical capabilities, consistent with findings in Andreadis et al. (2025). For regional policy, this finding indicates that non-tech hubs cannot rely on AI adoption alone to raise

wages. They must also simultaneously invest in reskilling and specialized educational infrastructure to enable the high-value skill recombination found in established tech hubs.

Limitations and Future Research Despite the robustness of our transformer-based metrics and overall results, this study is limited in the following ways. First, our analysis is restricted to publicly traded firms and job postings with explicitly listed salaries. This may over-represent established firms and overlook more volatile wage dynamics in startup and private equity-backed sectors. Second, while our Generative AI approach successfully identifies latent skills, it cannot account for skills that are omitted due to extreme commonality or obviousness in certain niches, potentially introducing residual measurement error in the estimated skill intensities. Third, we do not observe labor supply. The wage premia inferred from job postings, therefore, reflect contemporaneous demand signals and may not correspond to long-run equilibrium wages as the supply of AI-trained workers changes. Future research can leverage LinkedIn profiles to examine how workers are responding to labor market changes through reskilling. Finally, our model does not directly observe the actual use of AI within firms. We only observe the posted demand for AI skills, as reflected in job advertisements. Future work could link these skill-based measures to internal productivity data to determine whether these wage premia translate into realized firm-level productivity gains.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In Card, D. and Ashenfelter, O., editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

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

Acemoglu, D., Autor, D., and Johnson, S. (2023). Can we have pro-worker AI? Choosing a path of machines in service of minds. CEPR Policy Insight 123, CEPR Press.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44(2):139–160.

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour Economics*, 71:102002.

Andreadis, L., Kalotychou, E., Chatzikonstantinou, M., Louca, C., and Makridis, C. A. (2025). Local heterogeneity in artificial intelligence jobs over time and space. In *AEA Papers and Proceedings*, volume 115, pages 29–34. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Autor, D. and Salomons, A. (2018). Is automation labor-displacing? productivity growth, employment, and the labor share. Technical report, National Bureau of Economic Research.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Bessen, J. (2019). Automation and jobs: When technology boosts employment. *Economic Policy*, 34(100):589–626.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1):83–108.

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

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *The Econometrics Journal*, 14(3):368–386.

Dai, R., Liu, T., and Wang, C. R. (2025). Employee turnover in the spotlight: Evidence from mergers and acquisitions. Wharton research paper, Wharton Research Data Services (WRDS). Posted June 3, 2025; revised August 20, 2025.

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

Deming, D. and Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dixon, J., Hong, B., and Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Fedyk, A. and Hodson, J. (2023). Trading on talent: Human capital and firm performance. *Review of Finance*, 27(5):1659–1698.

Goldfarb, A., Taska, B., and Teodoridis, F. (2023). Could machine learning be a general purpose technology? evidence from online job postings. *Research Policy*, 52(1):104653.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Greenacre, M. (2021). Compositional data analysis. *Annual Review of Statistics and its Application*, 8(1):271–299.

Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.

Hemelt, S. W., Hershbein, B., Martin, S., and Stange, K. M. (2023). College majors and skills: Evidence from the universe of online job ads. *Labour Economics*, 85:102429.

Hershbein, B. and Kahn, L. B. (2016). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. NBER Working Paper 22762, National Bureau of Economic Research.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.

Huang, S., Seethor, B., Durmus, E., Handa, K., McCain, M., Stern, M., and Ganguli, D. (2025). How AI is transforming work at Anthropic. Available at: https://anthropic.com/research/how-ai-is-transforming-work-at-anthropic/.

Josten, C., Krause, H., Lordan, G., and Yeung, B. (2024). What skills pay more? the changing demand and return to skills for professional workers. Technical report, IZA Discussion Papers.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.

Lindquist, M. J. (2004). Capital–skill complementarity and inequality over the business cycle. *Review of Economic Dynamics*, 7(3):519–540.

Mäkelä, E. and Stephany, F. (2024). Complement or substitute? how ai increases the demand for human skills. *arXiv preprint arXiv:2412.19754*.

ONET Development (2025). Using a hybrid artificial intelligence-expert method to develop work style ratings for the O*NET database. Technical report, U.S. Department of Labor, Employment and Training Administration.

PwC AI Jobs Barometer Report (2025). Ai jobs barometer: Industry. https://pwc.turtl.co/story/ai-jobs-barometer-industry/page/6/1.

Tambe, P. B. (2025). Reskilling the workforce for ai: Domain expertise and algorithmic literacy. *Management Science*. Forthcoming.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Woessmann, L. (2024). Skills and earnings: A multidimensional perspective on human capital. *Annual Review of Economics*, 17.

Yin, W., Hay, J., and Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. *arXiv preprint arXiv:1909.00161*.

Appendix

A1 Additional Tables and Figures

Table A1: Lightcast Job Postings Data: Coverage and Linkage Summary

	Raw Lightcast Data		Compustat	IT Jobs	Salary
Year	Postings	Skills	Mapped	ONET15	Included
2018	33,125,641	417,279,417	9,459,777	1,115,564	87,263
2019	34,545,101	450,644,471	8,796,142	1,024,247	70,228
2020	35,072,775	435,835,184	8,528,923	852,387	53,469
2021	43,711,124	557,527,894	10,369,670	1,069,912	113,299
2022	48,727,337	652,293,853	10,986,588	1,323,175	235,335
2023	39,035,912	511,299,116	8,071,361	694,128	311,216
Total	234,217,890	3,291,410,613	56,212,461	6,079,413	870,810

Notes: This table presents the coverage of Lightcast job postings data from 2018–2023. Raw Lightcast Data includes all job postings and extracted skills. Compustat Mapped indicates successful linkage to public firm records. IT Jobs are defined by O*NET occupation code 15. Salary Included indicates verified salary information.

Figure A1: Cumulative Explanatory Power of Skill Domains

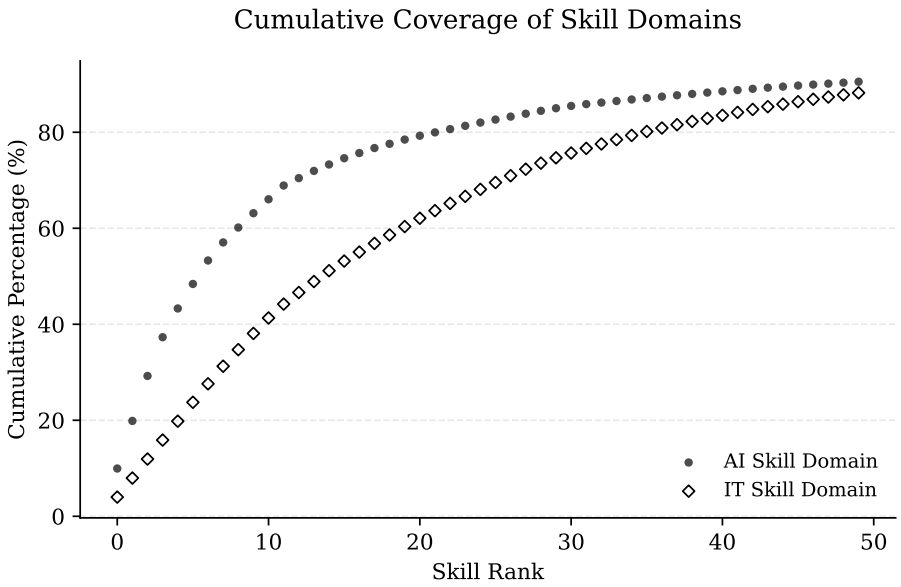


Table A2: Representativeness Check: Salary Disclosure vs. Full IT Universe

	Salary Sample		Full Universe		
	Mean	SD	Mean	SD	SMD
Panel A: Job Requirements					
Min Exp Required	5.19	3.08	5.26	3.01	-0.024
Education (Level 1-5)	1.82	0.68	1.86	0.63	-0.066
Skill Count	24.22	15.18	24.65	14.02	-0.031
Panel B: Firm Characteristics					
Log Total Assets ($\ln at$)	9.70	2.14	10.10	2.21	-0.178
Log Market Cap ($\ln mkvalt$)	9.71	2.01	10.07	2.04	-0.172
Panel C: Work Arrangement					
Remote: Not Stated	0.71	0.46	0.86	0.35	-0.150
Remote: On-site	0.03	0.16	0.01	0.10	0.015
Remote: Hybrid	0.03	0.16	0.01	0.10	0.014
Remote: Fully Remote	0.24	0.43	0.12	0.33	0.121
Panel D: Geography					
Transparency Law State	0.39	0.49	0.36	0.48	0.032
Total Observations: 870,810 (Salary) / 5,208,603 (Non-Salary)					

Notes: This table assesses the representativeness of the salary-disclosing sample by comparing it to the non-disclosing population of IT postings. Standardized Mean Difference (SMD) is calculated as $(\mu_{sal} - \mu_{full})/\sigma_{full}$. SMDs below 0.1 indicate statistical balance. Panel A covers human capital requirements, Panel B covers firm financial scale from Compustat, and Panels C-D cover work arrangements and regional characteristics.

Figure A2: Wage and Skill Trend (2018–2023)

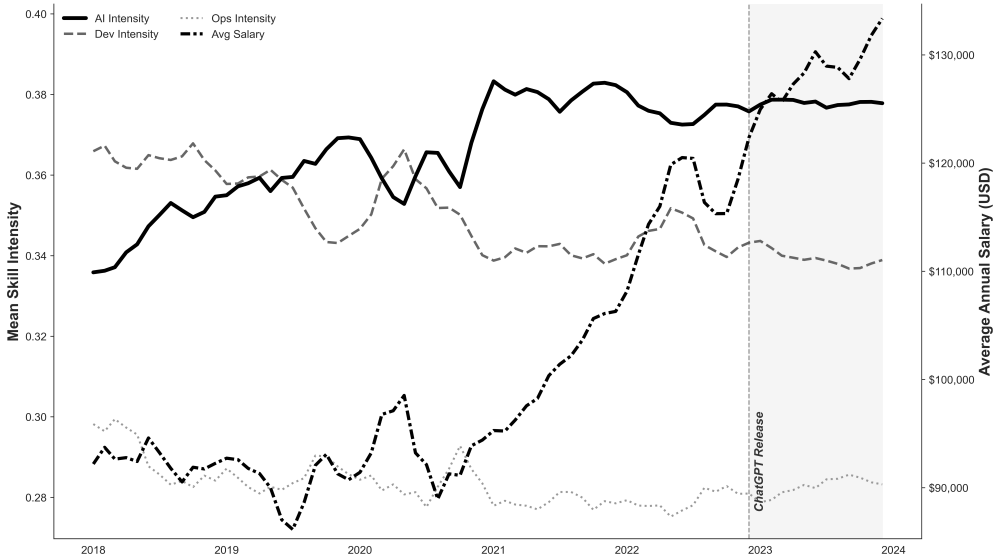


Table A3: Summary Statistics

	Mean	SD	Min	P25	P50	P75	Max
Panel A: Outcome Variable							
Salary ($'000s)	113.15	49.73	10.03	81.13	110.62	140.50	978.88
Log(Salary)	11.53	0.50	9.21	11.30	11.61	11.85	13.79
Panel B: Skill Scores							
IT Operations Skills	0.28	0.08	0.00	0.25	0.28	0.30	0.97
IT Development Skills	0.35	0.09	0.01	0.31	0.34	0.37	0.99
AI Skills	0.37	0.07	0.01	0.35	0.38	0.40	0.97
Panel C: Education							
Education: High School	0.06	0.24	0.00	0.00	0.00	0.00	1.00
Education: Associate	0.04	0.19	0.00	0.00	0.00	0.00	1.00
Education: Bachelor	0.58	0.49	0.00	0.00	1.00	1.00	1.00
Education: Master	0.20	0.40	0.00	0.00	0.00	0.00	1.00
Education: PhD	0.07	0.25	0.00	0.00	0.00	0.00	1.00
Education: Missing	0.34	0.48	0.00	0.00	0.00	1.00	1.00
Panel D: Experience							
Minimum Years Experience	3.67	3.50	0.00	0.00	3.00	5.00	15.00
Panel E: Work Arrangement							
Remote: On-site	0.73	0.44	0.00	0.00	1.00	1.00	1.00
Remote: Remote	0.24	0.43	0.00	0.00	0.00	0.00	1.00
Remote: Hybrid	0.03	0.16	0.00	0.00	0.00	0.00	1.00
Employment: Full-time	0.95	0.23	0.00	1.00	1.00	1.00	1.00
Employment: Part-time	0.03	0.18	0.00	0.00	0.00	0.00	1.00
Employment: Either	0.02	0.14	0.00	0.00	0.00	0.00	1.00
Internship Position	0.02	0.15	0.00	0.00	0.00	0.00	1.00
Number of Unique Firms: 2,709							
Total Observations: 870,797							

Notes: This table presents summary statistics for the main analysis sample. Panel A shows salary variables. Panel B reports skill scores aggregated into three categories: IT Operations (8 skills), IT Development (10 skills), and AI (12 skills). Panels C-E present covariates including education levels, work experience, and work arrangement characteristics. The remote category On-site is for postings that do not specify anything about work arrangement, and those that indicates postings that explicitly state that remote work is not allowed. The education categories do not add up to 1, as a job posting can demand one of multiple education levels. Regression also includes a missing indicator for the minimum years of experience. It also includes a measure of maximum years of experience (missing if not reported) for a very small fraction of postings that report this information.

Table A4: Out-of-Sample Prediction Performance

Year	N	MSE (Out-of-Sample)			R-squared (Out-of-Sample)		
		LLM Multi-dimensional Measure	Univariate AI Overlap Measure	% Reduction	LLM Multi-dimensional Measure	Univariate AI Overlap Measure	% Increase
Overall	256,257	0.195792	0.216268	9.47	0.1964	0.1123	74.81
2018	25,683	0.229918	0.272711	15.69	0.2765	0.1418	94.99
2019	20,608	0.225532	0.243067	7.21	0.2553	0.1974	29.33
2020	15,509	0.217456	0.243115	10.55	0.3023	0.2200	37.41
2021	32,736	0.194145	0.217703	10.82	0.2015	0.1046	92.64
2022	68,808	0.145895	0.157792	7.54	0.1437	0.0739	94.45
2023	90,726	0.152839	0.165852	7.85	0.1617	0.0903	79.07

Notes: Out-of-sample prediction results for log salary using 70-30 train-test split. Sample size corresponds to the 30% test data. LLM Multidimensional Measure denotes LASSO regression applied to centered log-ratio-transformed 30 skill scores. The AI Overlap Measure refers to the OLS regression with the AI overlap measure, following the steps provided in Babina et al. (2024) to construct the AI index. % Reduction denotes the percentage reduction in MSE for the LLM Multidimensional measure relative to the AI Overlap Measure. % Increase shows the percentage increase in R^2 for the LLM Multidimensional measure relative to the AI Overlap Measure.

Table A5: Out-of-Sample Prediction Performance: CLR vs. Ace(Narrow AI) vs. Ace(Broad AI)

Year	N	MSE (Out-of-Sample)			R-squared (Out-of-Sample)			
		CLR	Narrow	Broad	CLR	Narrow	Broad	% MSE Reduct (vs Broad)
2018–2023	257,585	0.196072	0.244079	0.239517	0.2025	0.0072	0.0258	18.14%
2018	25,797	0.229544	0.318357	0.315238	0.2811	0.0029	0.0127	27.18%
2019	20,730	0.225685	0.303265	0.299463	0.2588	0.0041	0.0166	24.64%
2020	15,702	0.216434	0.314012	0.309538	0.3152	0.0065	0.0207	30.08%
2021	33,084	0.194313	0.242187	0.239572	0.2018	0.0052	0.0159	18.89%
2022	69,074	0.146328	0.171961	0.170186	0.1514	0.0027	0.0130	14.02%
2023	90,960	0.153560	0.182367	0.179019	0.1658	0.0093	0.0275	14.22%

Notes: Out-of-sample prediction results for log salary using 70-30 train-test split. The out-of-sample prediction is on the test data-set (30% of the sample), consequently the sample sizes are lower than the full sample of job postings. CLR denotes LASSO regression applied to centered log-ratio-transformed 30 skill scores. Narrow AI Index follows Acemoglu et al. (2022) technical skill list. Broad AI Index adds general keywords "Artificial Intelligence" and "Machine Learning" to the Narrow list. % MSE Reduct denotes the percentage reduction in MSE for the CLR measure relative to the Broad AI Index.

Figure A3: Three-Dimensional Skills and Contribution to Estimated Wages

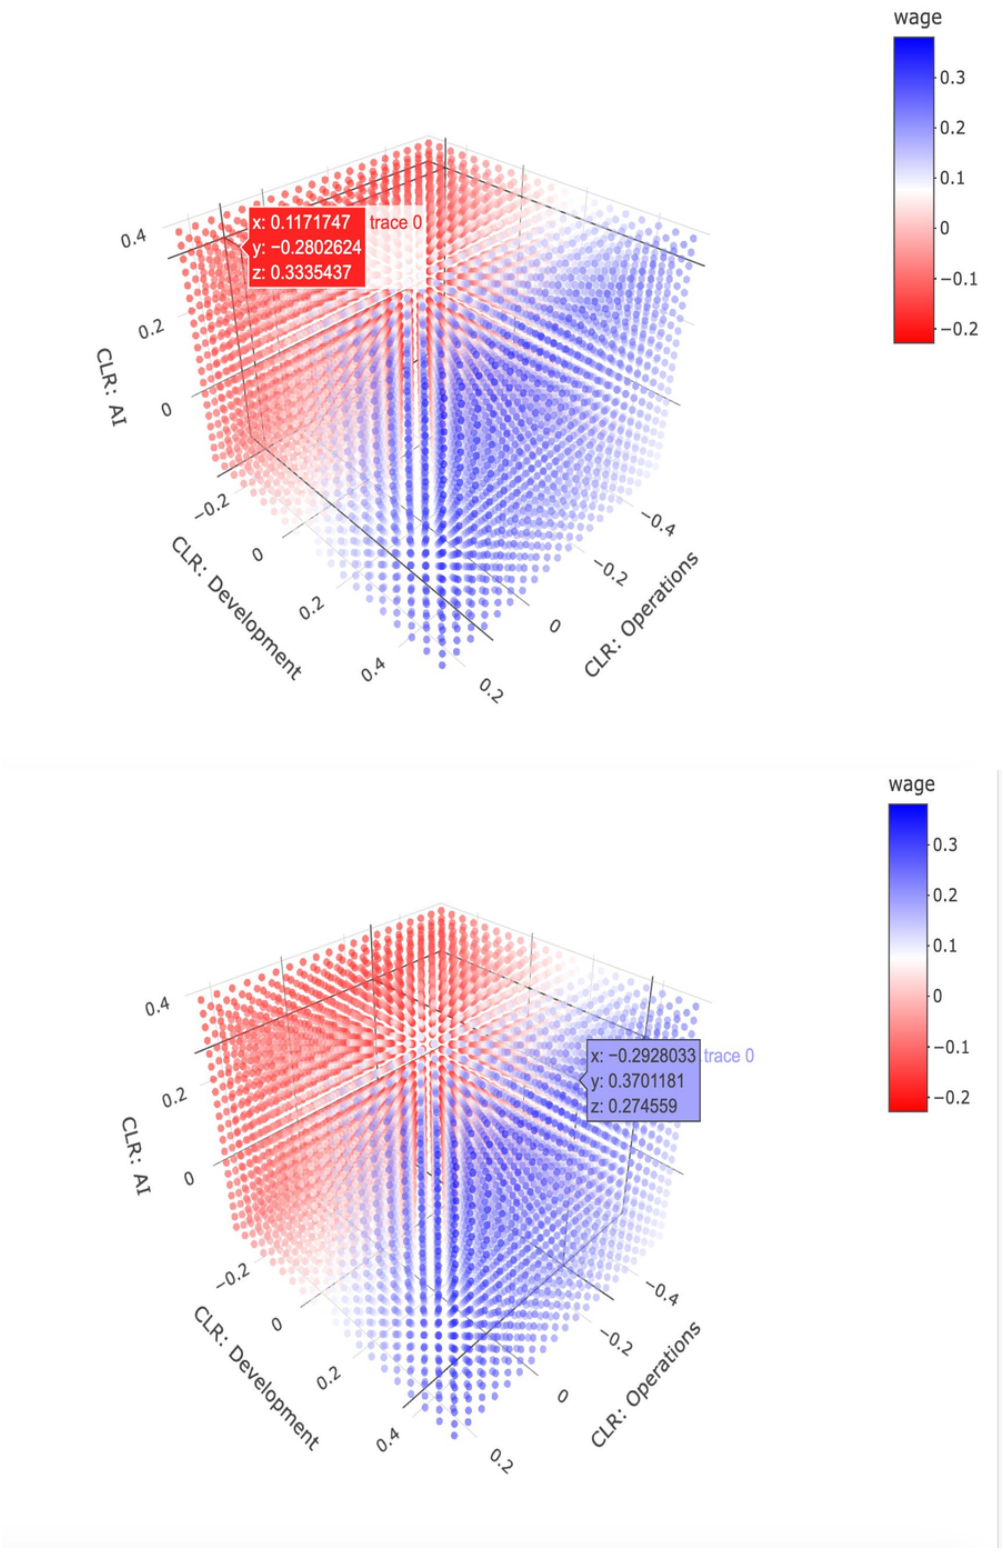


Table A6: Wage Premia using CLR Transformations (MSA Fixed Effects)

	Dependent variable: Log(Salary)		
	Main Effects	2-Way Interactions	3-Way Interactions
	(1)	(2)	(3)
CLR: Operations	0.050	0.160***	0.139***
	(0.059)	(0.043)	(0.045)
CLR: Development	0.315***	0.536***	0.530***
	(0.096)	(0.089)	(0.084)
Operations × Development		0.144***	0.200***
		(0.016)	(0.032)
Operations × AI		−0.293***	−0.231***
		(0.034)	(0.032)
Development × AI		0.226***	0.335***
		(0.055)	(0.058)
Operations × Development × AI			0.184***
			(0.054)
Firm FE	Yes	Yes	Yes
MSA FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	810,456	810,456	810,456
R^2	0.541	0.547	0.548
Adjusted R^2	0.539	0.545	0.546

Notes: *p<0.1; **p<0.05; ***p<0.01. Robustness check using MSA fixed effects instead of state fixed effects. Reference category: CLR AI (omitted). Controls include education, experience, remote work, employment type, and internship.

Table A7: Robustness: Controlling for Local AI Diffusion via Leave-Out MSA Exposure

	Dependent variable: Log(Salary)	
	MSA FE Only	MSA FE + AI Exposure
	(1)	(2)
CLR: Operations	0.134***	0.133***
	(0.044)	(0.043)
CLR: Development	0.524***	0.521***
	(0.085)	(0.084)
Operations × Development	0.202***	0.201***
	(0.031)	(0.031)
Operations × AI	-0.243***	-0.240***
	(0.030)	(0.030)
Development × AI	0.318***	0.319***
	(0.060)	(0.060)
Operations × Development × AI	0.185***	0.186***
	(0.053)	(0.053)
MSA-level AI Exposure (leave-out)		0.058***
		(0.011)
Firm FE	Yes	Yes
MSA FE	Yes	Yes
Year-Month FE	Yes	Yes
Controls	Yes	Yes
SE Clustered by Firm	Yes	Yes
Observations	789,333	789,333
R^2	0.535	0.536
Adjusted R^2	0.533	0.534

Notes: *p<0.1; **p<0.05; ***p<0.01. Column (2) controls for leave-out mean Babina AI exposure in the same MSA-month to account for local AI diffusion and labor market spillovers. The leave-out mean excludes the focal firm's own AI exposure to avoid mechanical correlation. Reference category: CLR AI (omitted). Standard errors clustered at firm level.

Table A8: Robustness Check: Quantile Regression

	Dependent variable: Log(Salary)		
	Q25	Q50	Q75
	(1)	(2)	(3)
CLR: Operations	0.142***	0.073***	0.053***
	(0.004)	(0.002)	(0.003)
CLR: Development	0.619***	0.475***	0.400***
	(0.004)	(0.003)	(0.004)
Operations × Development	0.254***	0.265***	0.231***
	(0.003)	(0.003)	(0.004)
Operations × AI	-0.195***	-0.182***	-0.157***
	(0.009)	(0.005)	(0.005)
Development × AI	0.521***	0.394***	0.315***
	(0.010)	(0.006)	(0.006)
Operations × Development × AI	0.286***	0.311***	0.244***
	(0.010)	(0.005)	(0.006)
Observations	870,797		
Firm + State + Year-Month FE	Yes (partialled out via Frisch-Waugh)		
Controls	Yes (partialled out)		

Notes: *p<0.1; **p<0.05; ***p<0.01. Columns report quantile regression estimates at the 25th, 50th, and 75th percentiles of the log salary distribution. Fixed effects and controls are partialled out prior to estimation via the Frisch-Waugh theorem (Canay, 2011). Reference category: CLR AI (omitted).

Table A9: Firm Sorting: Heterogeneity by Firm Age (Median Split)

	Dependent variable: Log(Salary)		
	Baseline	Below Median	Above Median
	(1)	(2)	(3)
CLR: Operations	0.139***	0.083	0.199***
	(0.049)	(0.056)	(0.021)
CLR: Development	0.543***	0.392***	0.662***
	(0.088)	(0.063)	(0.042)
Ops × Dev	0.205***	0.192***	0.178***
	(0.034)	(0.048)	(0.037)
Ops × AI	-0.235***	-0.205***	-0.195***
	(0.033)	(0.034)	(0.048)
Dev × AI	0.346***	0.274***	0.416***
	(0.060)	(0.063)	(0.094)
Ops × Dev × AI	0.185***	0.187***	0.116**
	(0.057)	(0.068)	(0.058)
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	805,927	426,976	378,951
Number of Firms	2,509	1,753	799
R^2	0.535	0.559	0.493
Adjusted R^2	0.533	0.557	0.492

Notes: *p<0.1; **p<0.05; ***p<0.01. Firm age is calculated as the posting year minus the founding year. Reference: CLR AI (omitted). Standard errors clustered at firm level.

Table A10: Robustness: Alternative Skill Aggregations

	Dependent variable: Log(Salary)		
	Baseline	Alt 1: DB→D	Alt 2: Cloud/DevOps→O
	(1)	(2)	(3)
CLR: Operations	0.137***	0.166***	0.130***
	(0.046)	(0.048)	(0.036)
CLR: Development	0.536***	0.598***	0.368***
	(0.085)	(0.087)	(0.082)
Ops × Dev	0.206***	0.195***	0.188***
	(0.033)	(0.032)	(0.029)
Ops × AI	-0.233***	-0.257***	-0.168***
	(0.032)	(0.029)	(0.043)
Dev × AI	0.342***	0.335***	0.240***
	(0.060)	(0.062)	(0.059)
Ops × Dev × AI	0.189***	0.160***	0.188***
	(0.054)	(0.051)	(0.048)
Operations Skills	8	7	10
Development Skills	10	11	8
AI Skills	12	12	12
Firm FE	Yes	Yes	Yes
State FE	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes
Observations	870,797	870,797	870,797
R^2	0.538	0.539	0.532
Adjusted R^2	0.537	0.537	0.530

Notes: *p<0.1; **p<0.05; ***p<0.01. Baseline: Operations (8), Development (10), AI (12). Alt 1: Database Management moved from Operations to Development. Alt 2: Cloud Computing and DevOps moved from Development to Operations. Reference: CLR AI (omitted). Standard errors clustered at the firm level.

Table A11: Robustness: Sample Restrictions

	Dependent variable: Log(Salary)			
	Baseline	Trim 1%	Firms≥5	Firms≥10
	(1)	(2)	(3)	(4)
CLR: Operations	0.137***	0.109**	0.136***	0.136***
	(0.046)	(0.043)	(0.046)	(0.047)
CLR: Development	0.536***	0.504***	0.536***	0.536***
	(0.085)	(0.083)	(0.085)	(0.085)
Ops × Dev	0.206***	0.206***	0.206***	0.207***
	(0.033)	(0.032)	(0.033)	(0.033)
Ops × AI	-0.233***	-0.197***	-0.233***	-0.233***
	(0.032)	(0.033)	(0.032)	(0.032)
Dev × AI	0.342***	0.364***	0.343***	0.344***
	(0.060)	(0.055)	(0.060)	(0.060)
Ops × Dev × AI	0.189***	0.200***	0.190***	0.191***
	(0.054)	(0.052)	(0.054)	(0.054)
Firm FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes	Yes
Observations	870,797	854,458	869,024	866,542
Number of Firms	2,709	2,693	1,848	1,478
R^2	0.538	0.560	0.537	0.537
Adjusted R^2	0.537	0.559	0.536	0.536

Notes: *p<0.1; **p<0.05; ***p<0.01. Column (1) replicates baseline specification. Column (2) drops top and bottom 1% of salary distribution. Columns (3) and (4) restrict to firms with at least 5 and 10 postings, respectively. Reference: CLR AI (omitted). Standard errors clustered at firm level.

Table A12: Mean Non-Technical Skill Counts and Standard Errors by Experience Level

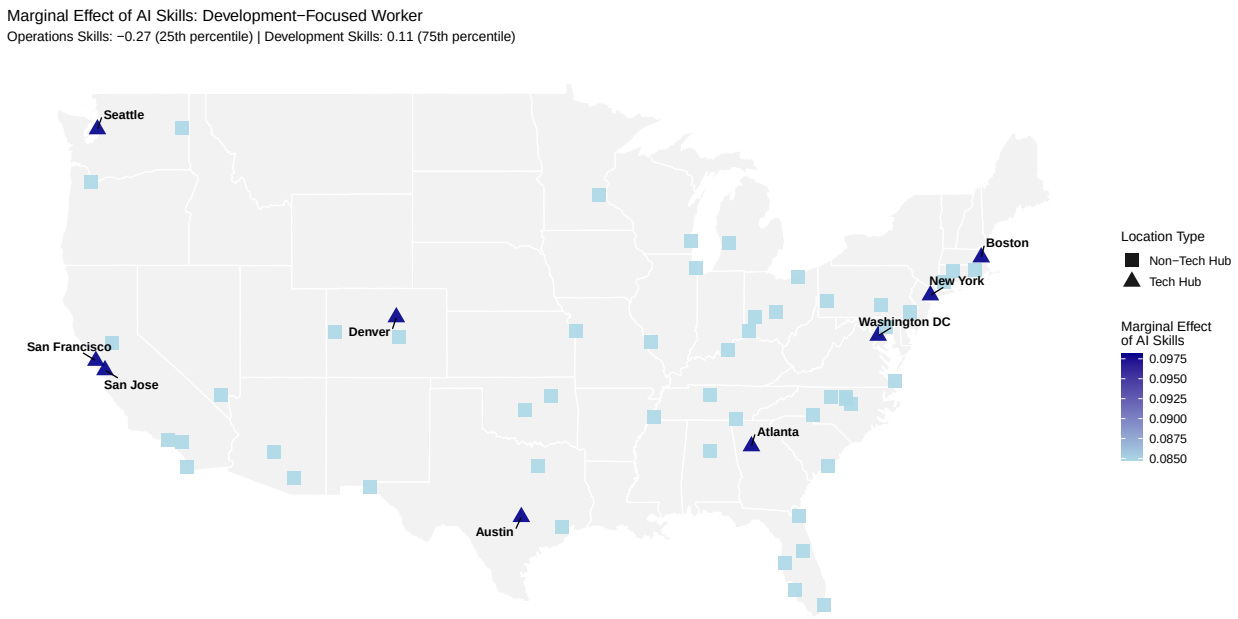
	Experience Level		
Variable	Entry	Mid	Senior
People Management	0.365	0.616	0.982
	(0.0013)	(0.0016)	(0.0032)
Project Management	0.034	0.053	0.071
	(0.0003)	(0.0004)	(0.0007)
Cognitive	0.910	1.128	1.323
	(0.0018)	(0.0021)	(0.0037)
Social	0.697	0.936	1.115
	(0.0016)	(0.0019)	(0.0033)
Observations	367,750	369,784	133,263

Notes: Table reports mean number of distinct non-technical skill tags per job posting matched within each category, with standard errors in parentheses. Skill categories follow Deming and Kahn (2018): People Management includes leadership, mentorship, coaching, and team building; Cognitive includes analytical thinking, problem solving, and strategic thinking; Social includes communication, collaboration, and relationship management; Project Management includes agile, scheduling, and project planning. Experience tiers are defined as Entry (0–2 years), Mid (3–7 years), and Senior (8+ years).

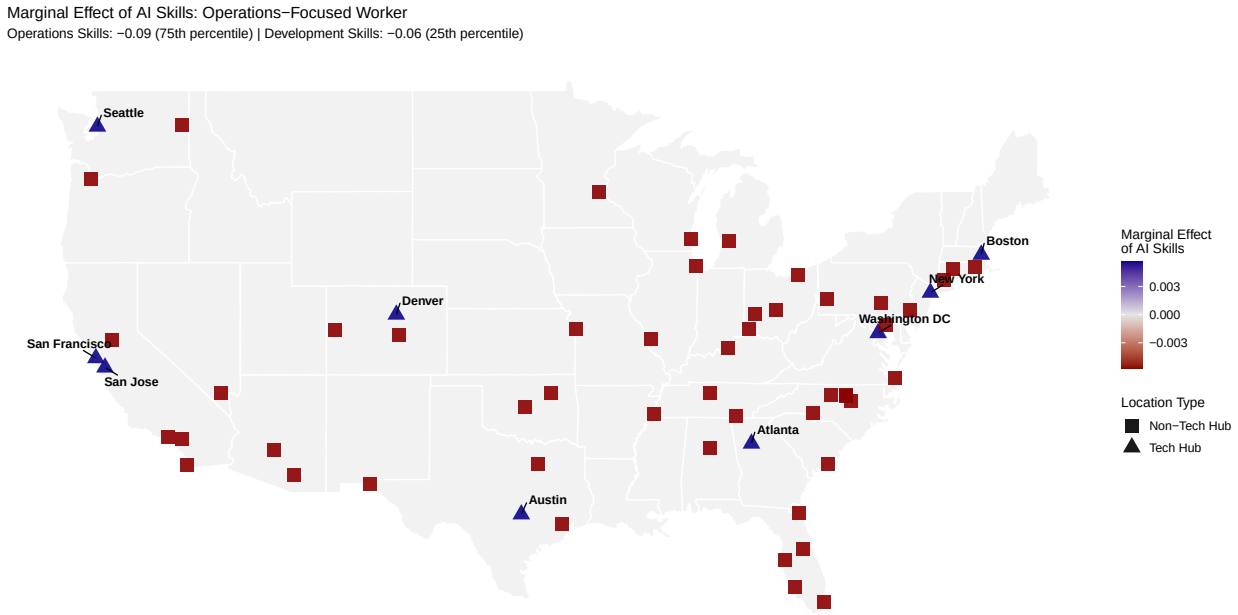
Table A13: Wage Premia: Tech Hubs vs Non-Tech Hubs

	Dependent variable: Log(Salary)	
	Tech Hub	Non-Tech Hub
	(1)	(2)
CLR: Operations	0.109*	0.140***
	(0.056)	(0.037)
CLR: Development	0.492***	0.539***
	(0.104)	(0.076)
Operations × Development	0.213***	0.195***
	(0.044)	(0.036)
Operations × AI	−0.257***	−0.176***
	(0.047)	(0.031)
Development × AI	0.322***	0.387***
	(0.080)	(0.058)
Operations × Development × AI	0.240***	0.176***
	(0.059)	(0.059)
Firm FE	Yes	Yes
State FE	Yes	Yes
Year-Month FE	Yes	Yes
Controls	Yes	Yes
SE Clustered by Firm	Yes	Yes
Observations	267,367	543,089
R^2	0.478	0.560
Adjusted R^2	0.475	0.558

Figure A4: Marginal Effect of AI and Tech Hubs



a. Development Focus



b. Operations Focus

Notes: This graph displays the marginal effects of AI on wages for two worker profiles where panel (a) is for a Development heavy skill portfolio and panel (b) is for a Operations heavy skill portfolio. Here we only display a subsample of MSAs on the graph for clarity.

Table A14: Wage Premia by Location and Work Flexibility: 2×2 Analysis

	Dependent variable: Log(Salary)			
	TH + On-site	TH + Remote/Hybrid	NTH + On-site	NTH + Remote/Hybrid
	(1)	(2)	(3)	(4)
CLR: Operations	0.135**	−0.025	0.180***	−0.028
	(0.056)	(0.029)	(0.033)	(0.018)
CLR: Development	0.539***	0.242***	0.585***	0.304***
	(0.106)	(0.044)	(0.068)	(0.042)
Operations × Development	0.214***	0.177***	0.179***	0.236***
	(0.046)	(0.030)	(0.036)	(0.026)
Operations × AI	−0.288***	−0.128***	−0.186***	−0.141***
	(0.051)	(0.024)	(0.030)	(0.028)
Development × AI	0.332***	0.227***	0.378***	0.305***
	(0.085)	(0.031)	(0.067)	(0.050)
Operations × Development × AI	0.226***	0.223***	0.121*	0.309***
	(0.064)	(0.033)	(0.062)	(0.039)
Firm FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
SE Clustered by Firm	Yes	Yes	Yes	Yes
Observations	206,446	60,921	396,505	146,584
R^2	0.482	0.524	0.555	0.618
Adjusted R^2	0.477	0.515	0.552	0.614

Notes: *p<0.1; **p<0.05; ***p<0.01. TH = Tech Hub, NTH = Non-Tech Hub. Tech Hubs include: San Jose, San Francisco, Seattle, Austin, Boston, Denver, Atlanta, NYC, and DC. Reference: CLR AI (omitted). Standard errors clustered at the firm level. The sample size for this analysis does not contain about 6.9% of the observations from the primary analysis sample due to missing MSA information on the job posting location.