The Great Reversal: Generative AI, Firm Productivity, and the Value of Specialized Talent*

Arzu Isik and Yunus Topbas

March 26, 2026

Abstract

This paper examines how artificial intelligence (AI) adoption by employees affects firm productivity. We first construct a novel global firm-year panel (2015–2025) by merging granular worker-level data with standardized financials. We then employ fixed-effects and instrumental variables to isolate causal effects, finding that broad AI adoption yields a 5.74% full-sample increase in revenue per employee per percentage-point increase in upskilled workers, with gains concentrated among large, listed firms in high-income countries. Smaller and upper-middle-income country firms experienced limited or negative returns. The advent of generative AI (e.g., ChatGPT) triggered a structural reversal. This shock severely degraded the baseline premium, reducing the pre-2023 broad adoption productivity gain from 10.52% down to 5.17% post-release. Commoditization of basic AI skills penalized smaller and private firms. In contrast, returns to dedicated AI specialists—deep-tech engineers focused on explicit AI software production or proprietary AI integration—surged, capturing a 15.60% global premium post-shock. These results indicate that as off-the-shelf AI becomes ubiquitous, value creation shifts from broad adoption toward firms with the structural scale and specialized human capital required for deep architectural integration.

Keywords: Artificial Intelligence (AI), Firm Productivity, Generative AI, AI Adoption, AI Specialization

*Arzu Isik, University of Calgary and Harvard University, arzu.isiktopbas@ucalgary.ca; Yunus Topbas, Peking University HSBC Business School, ytopbas@phbs.pku.edu.cn.

1 Introduction

The integration of artificial intelligence (AI) represents one of the most profound technological opportunities for enhancing modern firm productivity. While a burgeoning literature identifies AI as a critical form of intangible capital, the precise mechanisms through which this capital translates into measurable firm productivity remain opaque. In particular, does the broad diffusion of basic AI skills across a traditional workforce genuinely drive productivity, or does it merely represent superficial upskilling? Furthermore, who are the true beneficiaries of this technological diffusion, and how does a paradigm-altering leap—such as the advent of Generative AI—reallocate the "winners" and "losers" of the modern economy?

This paper bridges these empirical gaps by constructing a novel, globally comprehensive firm-year panel that links granular worker-level microdata from Revelio Labs with standardized cross-country financial statements from Bureau van Dijk's Orbis database. Our primary empirical focus is on evaluating the economic value of *AI Adoption*—the widespread operational upskilling of traditional workers (e.g., financial analysts or marketing professionals utilizing supplementary machine learning tools).

Utilizing a robust fixed-effects framework and an instrumental variable (IV) strategy to address endogeneity, we document that the productivity gains from broad AI adoption depend critically on a firm's preexisting scale and the macroeconomic environment. To accurately isolate the effect of AI adoption, we control for the presence of *AI Specialists* (e.g., dedicated "AI Engineers"). Over the full sample period, baseline broad AI adoption yielded positive and statistically significant returns, corresponding to a 5.74% global increase in revenue per employee for every 1 percentage point increase in the share of broadly upskilled workers. However, these gains were heavily skewed toward firms with significant preexisting resources. The clear baseline "winners" were predominantly large, publicly listed enterprises and firms domiciled in High-Income Countries (HIC). In contrast, smaller firms and those in Upper-Middle-Income Countries (UMC) exhibited no statistically significant gains, and in some scale classifications suffered immediate penalties from attempting broad adoption, suggesting that they might not have required infrastructure.

We further analyze the late-2022 release of Large Language Models (e.g., ChatGPT) to understand how a major exogenous technological shock alters the value of this intangible capital. Our evidence highlights a profound structural reversal: the democratization and

widespread availability of Generative AI effectively commoditized basic AI skills, eroding the competitive advantage of general workforce upskilling. Globally, this shock severely degraded the baseline premium, reducing the 10.52% global increase in productivity per percentage point of broad AI adoption down to 5.17% following the release of ChatGPT. The primary "losers" in this new generative regime were smaller enterprises, private firms, and Upper-Middle-Income Countries. For these groups, earlier investments in superficial upskilling were rendered obsolete by off-the-shelf automated solutions, triggering severe net productivity penalties (e.g., a 26.20% downward shock that dragged the net return to a negative 3.38% in UMCs). Conversely, large and publicly traded firms largely weathered the shock, utilizing their scale and governance to successfully insulate themselves from this commoditization.

Furthermore, our analysis yields a novel finding regarding the distinct trajectory of dedicated AI specialists. Crucially, because AI Adopters are distributed across a wide array of traditional corporate functions while AI Specialists are overwhelmingly concentrated in deep-tech software production, we do not directly benchmark their returns against one another. Instead, we isolate the intensive margin of explicit AI software production to cleanly identify who benefits from general AI diffusion, while simultaneously uncovering the independent economic trajectory of deep-tech specialization. While the generative boom eroded the value of general AI upskilling, it significantly amplified the returns to explicit, deep-tech role formalization. Firms with the organizational capacity to fundamentally integrate complex models into their proprietary architecture captured a highly significant 15.60% global premium on specialized talent post-shock. Temporal decompositions confirm that these dual macroeconomic forces—the accelerating penalty for broad adoption and the expanding premium for specialized expertise—intensified consistently through 2024 and 2025.

These findings remain highly consistent when restricting our analysis to knowledge-intensive, AI-exposed industries, and when utilizing raw, unweighted headcounts to validate our sampling methodologies. Ultimately, our results suggest that navigating the intangible transition of AI requires far more than the superficial adoption of accessible software; as basic skills commoditize, the locus of value creation shifts precipitously toward firms with the structural scale and specialized human capital required for deep architectural integration.

Our analysis contributes to several intersecting strands of literature: the macroeconomic implications of artificial intelligence, firm-level productivity, labor market dynamics, and the

role of intangible capital. Foundational scholarship views AI as a general-purpose technology capable of catalyzing growth or automating labor (Agrawal et al., 2019; Brynjolfsson, 2022), though Acemoglu (2025) cautions that reliably automating complex tasks remains difficult. We shift this macroeconomic focus to the firm, examining the organizational channels through which AI translates into productivity.

Our work also contributes to the empirical literature linking AI to firm performance. While AI drives growth through product innovation Babina et al. (2024), alters systematic risk (Babina et al., 2023), and generates valuation premiums for firms with established infrastructure Eisfeldt et al. (2023), much of this work abstracts from implementation heterogeneity. Furthermore, Barrios et al. (2025) highlight the prevalence of superficial "AI washing". We extend these studies by demonstrating that true productivity gains stem from formally embedding specialized AI human capital, rather than mere signaling or broad workforce upskilling.

Finally, we contribute to research on labor market adjustments and organizational capital. Transitioning from early concerns about job displacement (Acemoglu & Restrepo, 2020) to task reallocation (Hampole et al., 2025), we show that AI's productivity effects depend heavily on organizational strategies. Historical evidence proves technological diffusion requires substantial co-investment in organizational structure (Bloom et al., 2012; Brynjolfsson et al., 2021). Viewing generative AI as a structural shock that expands automatable tasks (Eloundou et al., 2023) and equalizes lower-skilled workers (Noy & Zhang, 2023), we frame AI as an organizational shock and a form of intangible capital. Ultimately, basic AI skills are commoditized by this equalization effect, making structural specialization and deep technical expertise the true drivers of productivity premiums.

2 Data and Stylized Facts

To investigate the allocative consequences and productivity premiums of AI human capital, we construct a novel, globally comprehensive firm-year panel. This requires linking granular, worker-level microdata with standardized, cross-country financial statements.

2.1 Data Sources and Panel Construction

Our foundational dataset relies on employment microdata provided by Revelio Labs, which aggregates and standardizes hundreds of millions of public employment profiles, resumes, and job postings globally. The raw data takes the form of a Job Spell Panel, where each observation corresponds to a unique position held by an individual, bounded by precise start and end dates. Crucially, Revelio's natural language processing algorithms parse the unstructured text of these profiles to extract specific technological skills and map job titles to standardized taxonomies.

We construct our primary firm-level measures of AI integration through a two-step aggregation process. First, we expand the individual Job Spell Panel into annual observations and aggregate them to an intermediate Firm-Occupation-Year level, utilizing the first five digits of `O*Net Code` to define occupation (comprising approximately 460 official SOC Broad Occupations). This intermediate step is critical for tracking historical occupational shares, which are subsequently used to construct our Shift-Share (Bartik) instrumental variables. Second, we collapse this panel to the Firm-Year level. To correct for the inherent sampling biases of online professional profiles, all firm-level aggregates are computed using Revelio's probability-of-representation sampling weights, ensuring our measures reflect the true structural composition of the firm's workforce.

We define two distinct variables to capture the depth and organizational design of AI integration. `AI Adoption` is defined as the sampling-weighted fraction of the firm's total workforce possessing general AI skills, regardless of their formal job title. This proxies for broad operational upskilling and the diffusion of AI tools across traditional workflows. `AI Specialist` is defined as the sampling-weighted fraction of the workforce holding formal, dedicated AI job titles, such as "Machine Learning Engineer" or "Director of AI." This proxies for explicit structural reorganization and the allocation of dedicated headcount to AI-centric functions.

We merge this aggregated employment panel with financial data from Bureau van Dijk's Orbis database to obtain our primary dependent variable, Revenue per Employee, alongside standard firm-level controls including total assets and firm age. Our matching method is outlined in Appendix B.

2.2 Sample Selection and Global Stratification

Given the scale of the global economy, analyzing the unconstrained population of firms introduces significant economic noise, as a vast majority of entities operate in sectors with negligible structural capacity for AI adoption, such as local retail or subsistence agriculture. To ensure our analysis focuses on economically meaningful entities, our primary filter is the successful intersection of granular worker-level microdata from Revelio Labs with standardized cross-country financial statements from the Orbis database. Crucially, across all empirical specifications, we impose a strict minimum threshold of 50 employees. This constraint filters out shell companies and ensure that the productivity metrics derived from Orbis revenue data are robust and representative of established organizational structures. This process yields a reliable baseline sample of firms with non-missing financial and occupational metrics.

To provide a comprehensive mapping of the returns to artificial intelligence, we stratify this global sample along three critical dimensions of organizational and macroeconomic heterogeneity. First, to evaluate international variance across stages of economic development, we map firms to World Bank income classifications based on their headquarters location. This allows us to partition our analysis into three distinct sub-samples: High-Income Countries (HICs), Upper-Middle-Income Countries (UMCs), and Lower-Middle-Income Countries (LMCs). Second, to account for the role of economies of scale in technological absorption, we categorize firms by their total workforce size into four groups: Small, Medium, Large, and Very Large. This stratification allows us to identify how preexisting organizational resources influence a firm's ability to extract value from broad AI upskilling. Finally, we distinguish between publicly listed and private firms to investigate the role of corporate governance and access to capital markets. By evaluating these intersecting layers of stratification, we can cleanly identify the specific organizational profiles that emerge as the winners and losers of the Generative AI transition.

2.3 Stylized Facts: The Evolution of AI Human Capital

Before turning to the formal analysis, we document several stylized facts regarding the demographics, occupational sorting, and evolutionary trajectory of the AI workforce, with figures presented in Appendix C. We partition our sample period into three distinct technological

eras: Episode 1 (2008–2014, the early machine learning era), Episode 2 (2015–2022, the deep learning boom), and Episode 3 (2023–2025, the Generative AI shock).

First, the demographic composition of AI talent is characterized by persistent gender disparities and high educational barriers to entry. Across our entire sample period, the predicted gender breakdown remains strikingly static, with males constituting approximately 70% of the AI workforce, a ratio that shows virtually no mean reversion even in the post-2023 era. Furthermore, AI adoption is fundamentally constrained by formal human capital accumulation. Over the three episodes, we observe a distinct rightward shift in the distribution of educational attainment. While Bachelor's degrees were the modal qualification during Episode 1, Master's degrees overtook them during Episode 2 and further solidified their dominance in Episode 3. Consistent with this, an overwhelming and growing majority of AI-exposed positions are classified under the "Extensive" preparation category within the O*NET Job Zone taxonomy, highlighting the deeply specialized, non-substitutable nature of this intangible capital.

Second, occupational sorting reveals a heavy concentration in technical domains. "Software Engineer" constitutes the single largest role category, and its relative share has expanded consistently across all three episodes. General "Engineer" and "Project and IT Specialist" roles represent the next largest cohorts. Non-technical functions—such as Sales, Marketing, and Operations—comprise only a marginal fraction of the direct AI workforce, reinforcing the notion that AI integration during this period remained a fundamentally technical, infrastructure-heavy endeavor.

Third, analyzing keyword adoption at the micro-level exposes the rapid democratization and shifting complexity of AI technologies. When classifying AI keywords by technical depth—such as Basic (e.g., "Machine Learning," "Prompt Engineering"), Intermediate (e.g., "Natural Language Processing," "LLM"), and Advanced (e.g., "Convolutional Neural Network," "PyTorch")—we observe a secular increase in the fraction of "Basic" AI keywords relative to "Intermediate" and "Advanced" terms across the three episodes. This suggests a broadening of basic AI literacy across the general workforce as enterprise tools become more user-friendly. However, text analyses of the most prevalent technical keywords capture the dramatic structural breaks defining our three episodes. During Episode 1, foundational terms such as "Machine Learning," "Neural Network," and "Computer Vision" dominated the landscape. Episode 2 witnessed the rise of specialized deep learning architectures and frame-

works, highlighted by the sudden prevalence of "Deep Learning," "TensorFlow," "PyTorch," and "Convolutional Neural Network." The onset of Episode 3 marks a clear paradigm shift: terms like "GenAI," "Generative Pretrained Transformer," and "LLM" (Large Language Model) entirely disrupt the frequency distribution (see Appendix A1 for the complete keyword taxonomy).

Finally, examining macroeconomic penetration reveals a profound geographic agglomeration and a widening technological divide. The annualized density of AI positions exhibits extreme stratification when decomposed by World Bank Income Groups. High-Income Countries (HICs) and Upper-Middle-Income Countries (UMCs) trace steep, exponential adoption curves that inflect upward around the deep learning boom of 2015 and accelerate further into the Generative AI era. Conversely, Lower-Middle-Income (LMCs) and Low-Income Countries (LICs) exhibit stagnant, near-zero baseline growth, illustrating a systemic friction in technological transfer. This vast divergence strongly motivates our empirical strategy to estimate productivity premiums separately across HIC, UMC, and LMC subsamples. Geographically, this inequality is visibly clustered. Global heat maps show AI capital heavily agglomerated in North America, Western Europe, and India. Within the United States, spatial clustering is remarkably pronounced; AI density is intensely concentrated within traditional technology hubs in California, Washington, New York, and Massachusetts.

Together, these stylized facts confirm that AI represents a highly specialized, rapidly evolving, and heavily agglomerated form of intangible capital. The shift from predictive machine learning (Episodes 1 and 2) to generative architectures (Episode 3) presents a unique macroeconomic shock, setting the stage for our empirical investigation into how these distinct forms of technological adoption drive firm boundaries and global productivity.

3 Empirical Methodology

To rigorously evaluate the impact of artificial intelligence integration on firm-level productivity, and to disentangle the value of broad operational upskilling from specialized structured organizations, we employ a firm-year panel fixed-effects framework. We estimate our baseline and interacted specifications across the consolidated Global sample, as well as separately within the HIC, UMC, and LMC sub-samples, to test for heterogeneous returns across different stages of global economic development.

3.1 Distinguishing AI Adoption from AI Specialists

A critical innovation of this paper is isolating the economic difference between a firm investing in general human capital upskilling versus executing a formal structural specialization. To do this, we rely on two highly distinct constructs from the Revelio Labs dataset: AI Adoption and AI Specialist.

For AI Adoption (General Upskilling), we parse job resume descriptions within the Revelio data for the keywords listed in Appendix A. Descriptions containing at least one keyword are tagged as AI Adopters. This measure is generated by text-mining self-reported skill tags and resume bullet points, capturing the penetration of AI as a supplementary tool into traditional workflows (e.g., a "Financial Analyst" or "Software Developer" whose core job title has not changed, but who employs machine learning models to execute their duties).

For AI Specialist (Structural Specialization), jobs with "AI" explicitly included in their official titles are tagged as AI Specialists. Unlike skill keywords, this flag is generated by algorithmically mapping a worker's formal, standardized job title to explicit AI categories within the highly granular `role_k17000` taxonomy of Revelio (e.g., mapping titles directly to "AI Sales Executive," "AI Product Manager," "AI Consulting," or "Enterprise AI Architect"). This identifies the formal structural specialization of the firm to dedicate specific headcount entirely to AI-centric functions.

3.2 Baseline Productivity and Organizational Design

We begin by modeling firm productivity as a function of the firm's allocation of AI human capital. The baseline specification is structured as follows:

$$Y_{ft} = \beta_1 AI_Adoption_{ft} + \beta_2 AI_Specialist_{ft} + X_{ft}\Gamma + \delta_f + \theta_{st} + \varepsilon_{ft} \tag{1}$$

The dependent variable, Y_{ft}, is the raw revenue per employee for firm f in year t (expressed in millions of dollars), serving as our primary proxy for labor productivity. The two parameters of central interest are β_1 and β_2, which capture the distinct economic channels through which AI is integrated into the firm. β_1 measures the productivity premium associated with $AI_Adoption_{ft}$, defined as the sampling-weighted share of the firm's total

workforce possessing general AI skills (a ratio from 0 to 1). This variable reflects the broad diffusion of AI as an operational tool among incumbent workers (e.g., traditional financial analysts or marketers utilizing machine learning algorithms). Conversely, β_2 isolates the premium associated with $AI_Specialist_{ft}$, defined as the sampling-weighted share of the workforce holding dedicated, formal AI job titles (also bounded between 0 and 1). This coefficient identifies the returns to formal structural reorganization and the intensive margin of specialized technological capital.

To isolate the causal effect of AI from confounding firm characteristics, the vector X_{ft} incorporates a robust set of time-varying firm-level controls. We include Firm Size (the natural logarithm of total assets) to control for economies of scale, and Firm Age to account for lifecycle productivity dynamics. Crucially, we control for Capital Intensity (the natural logarithm of tangible assets) to ensure that our AI premiums are not mechanically capturing highly capitalized, asset-heavy firms. Finally, to distinguish the specific effect of AI from general digitization, we construct a measure of Overall Tech Intensity—the sampling-weighted share of the firm's workforce employed in traditional IT, software, or data roles who do not possess AI skills. We include firm fixed effects (δ_f) to absorb all time-invariant, unobserved firm heterogeneity, such as persistent managerial quality, geographic advantages, or overarching corporate culture. We also include Industry-by-Year fixed effects (θ_{st}) to control for time-varying macroeconomic shocks, shifting consumer demand, and sector-specific technology waves. Standard errors are clustered at the Industry level to account for within-firm serial correlation over time.

3.3 The Generative AI Shock

While the baseline model captures the average productivity premium over the full sample, the release of advanced Large Language Models (LLMs) in late 2022 represented a profound structural break in the accessibility and capability of artificial intelligence. To explicitly test how the advent of Generative AI altered the organizational returns to technology, we augment the baseline model by interacting our human capital measures with a post-GPT indicator:

$$
\begin{aligned}
Y_{ft} = \beta_1 AI_Adoption_{ft} + \beta_2 AI_Specialist_{ft} \\
+ \beta_3 (AI_Adoption_{ft} \times GPT_t) + \beta_4 (AI_Specialist_{ft} \times GPT_t) \\
+ X_{ft}\Gamma + \delta_f + \theta_{st} + \varepsilon_{ft}
\end{aligned} \tag{2}
$$

Here, GPT_t is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The standalone main effect of GPT_t is fully absorbed by the Industry-Year fixed effects (θ_{st}), as the shock constitutes a macroeconomic, time-varying event experienced uniformly within sectors.

The coefficients on the interaction terms, β_3 and β_4, formalize our test of the commoditization versus complementarity hypotheses. A negative β_3 would suggest that the highly accessible, off-the-shelf nature of Generative AI tools eroded the competitive advantage of possessing basic, broad AI adoption, effectively commoditizing low-level AI literacy. In contrast, a positive β_4 would indicate that the integration of these highly complex, enterprise-level models requires deep technical expertise, thereby amplifying the productivity premium for structurally dedicated AI specialists.

3.4 Identification Strategy: Dual Shift-Share (Bartik) IV

The primary threat to causal inference in our fixed-effects framework is endogeneity driven by reverse causality and omitted variable bias. Highly productive, cash-rich firms are mechanically more capable of bidding for and hiring expensive AI talent, which could positively bias ordinary least squares (OLS) estimates. Because our models include two distinct, endogenous regressors (*AI_Adoption* and *AI_Specialist*), the rank condition requires a system of at least two independent instrumental variables.

To overcome this, we employ a dual Shift-Share (Bartik) instrumental variable approach. This strategy exploits the pre-existing, historical occupational structure of the firm, combined with national-level technological transformation rates, to predict exogenous variations in AI human capital allocation. We define the firm's historical base period as 2010–2014, predating the explosion of both deep learning and Generative AI.

Instrument 1 (For AI Adoption): The first instrument predicts the general upskilling

of the firm's workforce based on its legacy labor composition:

$$Z_{1,ft} = \sum_{o} \omega_{fo,base} \times \Delta AI_Skill_National_{ot}$$

where $\omega_{fo,base}$ is the historical share of occupation o (defined by granular, 460-category of the first five digits of the O*Net Code) at firm f during the base period. We interact this fixed, historical share with $\Delta AI_Skill_National_{ot}$, representing the national, leave-one-out rate at which workers in occupation o globally acquire general AI skills over time.

Instrument 2 (For AI Specialists): The second instrument predicts the firm's propensity to formally restructure its organizational chart to include dedicated AI roles:

$$Z_{2,ft} = \sum_{o} \omega_{fo,base} \times \Delta AI_Title_National_{ot}$$

where $\Delta AI_Title_National_{ot}$ is the national rate at which workers in historical occupation o formally transition into explicitly re-titled AI specialist roles over time.

The economic logic underpinning this dual-IV strategy relies on the independent variation in national occupational trajectories. For example, a firm historically reliant on traditional Web Developers will receive a high Z_1 score (because Web Developers rapidly learn basic AI tools nationally) but a comparatively low Z_2 score (because Web Developers rarely face formal structural re-titling to "AI Engineer"). Conversely, a firm historically heavy in Statisticians or Data Scientists will receive a high score for both instruments. This independent, exogenous variation in the national evolution of different occupations allows us to successfully identify and disentangle the causal effects of broad adoption (β_1) versus structural reorganizational specialization (β_2).

4 AI Adoption, Productivity, and the Generative AI Shock

In this section, we present our primary empirical findings on the impact of artificial intelligence integration on firm-level productivity, focusing specifically on the structural shifts induced by the Generative AI (GPT) shock. To provide a comprehensive understanding of these dynamics, we decompose our results along three critical dimensions. First, we examine the heterogeneity of AI productivity effects across global economic development stages by

partitioning the sample into World Bank income groups. Second, we analyze how economies of scale influence these returns by evaluating the effects across different firm size classifications. Finally, we investigate the role of corporate governance and capital market access by comparing the productivity premiums between publicly listed and private firms.

We first outline the primary economic mechanisms driving our results. Our primary variable of interest is *AI Adoption*—the broad upskilling of the general workforce. Crucially, we do not directly benchmark the returns of AI Adoption against *AI Specialists*. As Figure 8 reveals, AI Adopters are distributed across a wide array of traditional corporate functions (e.g., finance, sales, project management), whereas AI Specialists are overwhelmingly concentrated in a single domain (over 60% are Software Engineers). Consequently, a direct economic comparison risks conflating general technological diffusion with sector-specific software production. Instead, we utilize AI Specialists primarily as a rigorous control for the direct, deep-tech production of AI within a firm, while analyzing its heterogeneity separately as a secondary mechanism.

We find that the productivity gains from AI depend critically on a firm's preexisting scale and the macroeconomic environment. Before the introduction of Generative AI, broad AI adoption yielded substantial returns, but the clear winners were predominantly large, publicly listed enterprises and firms in High-Income Countries (HIC). The release of ChatGPT and similar Large Language Models altered this distribution by making advanced AI tools widely accessible. This shock effectively commoditized basic AI skills, eroding the competitive advantage of general workforce upskilling. The primary losers in this new regime were smaller, private firms and those in Upper-Middle-Income Countries (UMC), where earlier investments in broad adoption became a drag on productivity. In contrast, large public firms largely weathered the shock. More importantly, while the generative boom diminished the value of general AI adoption, it significantly amplified the returns to hiring dedicated AI specialists. Firms that possessed the organizational capacity to deeply integrate these complex models into their own architecture emerged as the ultimate beneficiaries, seeing their productivity premiums widen substantially.

Consistent with our empirical strategy, the dependent variable throughout this section is the natural logarithm of firm productivity (revenue per employee in millions). Standard production function controls, including the log of capital intensity (fixed assets per employee) and the log of the total workforce (employees), are included in all specifications. Because the

dependent variable is log-transformed and our AI metrics are bounded between 0 and 1, the coefficients are straightforward to interpret: a coefficient of β implies a β percent change in productivity for every 1 percentage point (1%) increase in the respective AI workforce share.

4.1 Heterogeneity in AI Productivity Effects by Income Group

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	5.74***	8.64***	-0.15	44.36***
	(1.51)	(1.51)	(7.52)	(10.89)
AI Specialist	-9.96	-14.46*	100.56**	-66.17
	(8.27)	(8.29)	(49.24)	(56.30)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 1: **Baseline Productivity and AI Integration**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** is the sampling-weighted fraction of the firm's workforce possessing general AI skills. **AI Specialist** is the sampling-weighted fraction of the workforce holding formal, dedicated AI job titles. The sample is decomposed across World Bank Income Groups: High-Income Countries (HIC), Upper-Middle-Income Countries (UMC), and Lower-Middle-Income Countries (LMC). **Capital** is defined as the logarithm of the ratio of total fixed assets to the total number of employees. **Labor** is defined as the logarithm of the total number of employees. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls (firm size, proxied by total assets; firm age; and overall tech intensity), as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 1 evaluates the baseline impact of AI integration across the global sample and World Bank income groups over the entire sample period, prior to isolating the GPT shock. The

standard production controls align with traditional economic assumptions: capital intensity yields a significant positive elasticity of 0.10 globally, while the employee scale parameter is significantly negative at -0.33, consistent with decreasing marginal returns to labor.

Focusing on our primary variable of interest, the global baseline return to broad AI adoption is positive and statistically significant, indicating that a 1 percentage point increase in the share of broadly upskilled workers corresponds to a 5.74% increase in revenue per employee. This global premium is largely driven by significant productivity boosts of 8.64% in High-Income Countries and 44.36% in Lower-Middle-Income Countries. In contrast, Upper-Middle-Income Countries (UMC) don't exhibit statistically significant gains from broad adoption.

Table 2 introduces the Generative AI interaction terms, capturing the commoditization effect of the GPT shock. Globally, the interaction for broad AI adoption is significantly negative, effectively offsetting the baseline premium by imposing a 5.35% penalty per percentage point increase in adoption. This shock varies noticeably across economic development stages. In HIC markets, the interaction imposes a significant 4.19% penalty on broad adoption. The decline is most pronounced in UMCs, where basic adoption post-GPT is associated with a significant 26.20% productivity penalty for every 1 percentage point increase in the upskilled workforce. This pattern suggests that as emerging markets attempted to integrate AI without sufficient complementary digital infrastructure, the sudden availability of automated generative tools reduced the relative value of their earlier upskilling efforts.

The annual decomposition of the GPT shock (Table 3) shows that this commoditization trend strengthens over time. The significant penalty for AI adoption moves from 4.82% in 2023 to 6.27% in 2024, and this penalty intensifies further in 2025. It should be noted that our financial data for 2025 is incomplete due to a collection cutoff; consequently, while the directional trend remains highly robust, the exact coefficient magnitudes for 2025 appear as extreme outliers and are therefore not interpreted at face value.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.52***	12.29***	22.82***	53.09***
	(1.66)	(1.67)	(7.91)	(12.70)
AI Specialist	-21.62**	-27.51***	107.10**	-60.21
	(8.77)	(8.80)	(53.27)	(65.11)
AI Adopter × GPT	-5.35***	-4.19***	-26.20***	-6.42
	(0.78)	(0.80)	(3.21)	(4.03)
AI Specialist × GPT	15.60***	16.74***	13.31	-0.74
	(3.50)	(3.63)	(15.14)	(19.45)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 2: **Impact of ChatGPT on The Relationship Between AI Integration and Productivity**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 1. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills (AI Adoption × GPT) and the complementarity of deep expertise (AI Specialist × GPT). All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 1). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.73***	12.49***	23.38***	53.52***
	(1.66)	(1.68)	(7.93)	(13.21)
AI Specialist	-22.94***	-29.51***	107.26**	-60.04
	(8.82)	(8.84)	(53.87)	(68.84)
AI Adopter × 2023	-4.82***	-4.44***	-22.53***	-6.00
	(0.72)	(0.75)	(3.06)	(3.82)
AI Adopter × 2024	-6.27***	-3.72***	-33.42***	-7.29
	(1.12)	(1.17)	(4.46)	(5.13)
AI Adopter × 2025	-34.15***	7.33	-146.08	-43.82***
	(11.52)	(6.15)	(119.08)	(14.39)
AI Specialist × 2023	10.98***	11.36***	6.57	-1.94
	(3.20)	(3.34)	(13.74)	(17.36)
AI Specialist × 2024	24.44***	28.38***	25.13	1.69
	(5.41)	(5.72)	(21.01)	(25.83)
AI Specialist × 2025	154.55***	107.71***	44.47	-153.64***
	(41.62)	(30.73)	(139.80)	(54.19)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 3: **Temporal Dynamics of the Generative AI Shock**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 1. The interaction terms evaluate the year-by-year acceleration of the commoditization effect for broad skills and the complementarity effect for specialized expertise following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 1). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Meanwhile, our secondary control variable—the share of AI specialists—follows a different path. At baseline (Table 1), the global coefficient for specialists is negative but statistically insignificant, meaning that a high concentration of software engineers did not universally translate to higher revenue per employee before the generative AI boom. Following the GPT shock (Table 2), however, the global interaction for specialists increases significantly, yielding a 15.60% premium. Table 3 confirms that this premium for dedicated AI specialists grows considerably, accelerating consistently from 2023 through 2025. While basic AI literacy became widely accessible, the specialized engineering skills required to build and integrate proprietary generative architectures commanded a rising premium.

4.2 Heterogeneity in AI Productivity Effects by Firm Size

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	5.74***	-4.47**	9.90***	14.35***	20.36***
	(1.51)	(2.04)	(2.73)	(4.52)	(4.63)
AI Specialist	-9.96	6.89	-2.22	-57.84**	-47.13
	(8.27)	(14.76)	(11.80)	(26.99)	(30.40)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.946	0.938	0.943	0.957

Table 4: **Baseline Productivity and AI Integration by Firm Size**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by firm size. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** is the sampling-weighted fraction of the firm's workforce possessing general AI skills. **AI Specialist** is the sampling-weighted fraction of the workforce holding formal, dedicated AI job titles. The sample is decomposed by firm size categories: Small, Medium, Large, and Very Large. **Capital** is defined as the logarithm of the ratio of total fixed assets to the total number of employees. **Labor** is defined as the logarithm of the total number of employees. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls (firm size, proxied by total assets; firm age; and overall tech intensity), as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

To understand how firm size influences the ability to absorb AI capital, we partition the global sample into small, medium, large, and very large firms (Tables 4, 5, and 6). The baseline results in Table 4 show a positive relationship between firm size and the returns to broad AI adoption. The estimated adoption effect shifts from a significant 4.47% penalty for small firms to positive, significant premiums of 9.90% for medium, 14.35% for large, and 20.36% for very large firms. Prior to Generative AI, extracting productivity gains from broad AI literacy generally required complementary investments in proprietary data, enterprise software, and computing power, which naturally favored larger organizations. Small firms attempting to upskill their workforce without this infrastructure saw negative returns.

The Generative AI shock (Table 5) alters this dynamic, placing a heavier burden on smaller firms. Post-GPT, the penalty for basic AI adoption is universally negative, with small firms experiencing the largest significant interaction penalty of 10.47%. Medium firms also face a significant 3.88% penalty. In contrast, large and very large firms are relatively insulated; their adoption interaction terms are negative but not statistically significant. The widespread availability of AI tools appears to have erased the competitive edge that smaller firms sought through basic upskilling, whereas larger firms with the scale to internalize these tools managed the transition more effectively.

Looking at the specialist control variable, the Generative AI shock highlights a complementarity effect that is concentrated among the largest firms. Although specialists provided ambiguous baseline returns, the post-GPT interaction (Table 5) for very large firms yields a significant 49.20% premium. The temporal breakdown (Table 6) shows this premium expanding substantially through 2025 (subject to the aforementioned 2025 data cutoff). Overall, the generative era reduced the value of broad adoption for smaller firms while disproportionately rewarding the largest enterprises equipped to deploy specialized, deep-tech teams.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.52***	4.36**	13.13***	17.21***	23.36***
	(1.66)	(2.18)	(3.08)	(5.07)	(5.21)
AI Specialist	-21.62**	-3.86	-1.42	-68.62**	-110.67***
	(8.77)	(15.43)	(12.69)	(29.20)	(32.91)
AI Adopter × GPT	-5.35***	-10.47***	-3.88**	-2.79	-1.80
	(0.78)	(0.91)	(1.55)	(2.26)	(2.06)
AI Specialist × GPT	15.60***	18.01***	1.34	11.56	49.20***
	(3.50)	(4.68)	(5.53)	(9.40)	(14.63)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.947	0.938	0.943	0.958

Table 5: **Impact of ChatGPT on The Relationship Between AI Integration and Productivity by Firm Size**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across different firm size categories. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 4. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills (AI Adoption × GPT) and the complementarity of deep expertise (AI Specialist × GPT). All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 4). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.73***	4.66**	13.23***	17.43***	23.43***
	(1.66)	(2.18)	(3.08)	(5.09)	(5.21)
AI Specialist	-22.94***	-5.09	-1.54	-72.34**	-113.75***
	(8.82)	(15.40)	(12.75)	(29.66)	(32.36)
AI Adopter × 2023	-4.82***	-8.34***	-3.97***	-3.29	-3.89**
	(0.72)	(0.89)	(1.43)	(2.21)	(1.98)
AI Adopter × 2024	-6.27***	-15.08***	-3.38	-1.70	1.02
	(1.12)	(1.31)	(2.30)	(3.15)	(2.43)
AI Adopter × 2025	-34.15***	-113.43***	-131.88***	-24.76	0.28
	(11.52)	(37.78)	(35.23)	(29.01)	(5.97)
AI Specialist × 2023	10.98***	12.84***	1.40	7.77	44.29***
	(3.20)	(4.62)	(5.05)	(9.32)	(12.98)
AI Specialist × 2024	24.44***	28.54***	0.76	17.84	48.81***
	(5.41)	(8.23)	(9.17)	(12.47)	(17.23)
AI Specialist × 2025	154.55***	150.90	705.24***	126.84	135.18***
	(41.62)	(112.78)	(272.49)	(79.65)	(40.53)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.947	0.938	0.943	0.958

Table 6: **Temporal Dynamics of the Generative AI Shock by Firm Size**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns across firm size classifications. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 4. The interaction terms evaluate the year-by-year acceleration of the commoditization effect for broad skills and the complementarity effect for specialized expertise following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 4). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.3 Heterogeneity in AI Productivity Effects by Listing Status

Finally, we investigate how corporate governance and access to capital markets shape technology diffusion by splitting the sample into publicly listed and private firms (Tables 7, 8, and 9). At baseline (Table 7), public firms realized productivity gains from broad AI adoption, seeing a 4.98% premium that is significant at 10% but not at 5%. This aligns with our firm-size findings: publicly traded entities typically have the capital access and organizational bandwidth necessary to finance the infrastructure required for technology integration. Private firms, which may face tighter capital constraints, showed an insignificant 1.87% return from AI adoption prior to the generative shock.

	Firm Productivity		
	Global	Public	Private
AI Adopter	5.74***	4.98*	1.87
	(1.51)	(2.82)	(1.48)
AI Specialist	-9.96	-7.02	-11.87
	(8.27)	(15.34)	(8.18)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 7: **Baseline Productivity and AI Integration by Listing Status**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by listing status. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** is the sampling-weighted fraction of the firm's workforce possessing general AI skills. **AI Specialist** is the sampling-weighted fraction of the workforce holding formal, dedicated AI job titles. The sample is decomposed into Public and Private firms. **Capital** is defined as the logarithm of the ratio of total fixed assets to the total number of employees. **Labor** is defined as the logarithm of the total number of employees. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls (firm size, proxied by total assets; firm age; and overall tech intensity), as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The transition following the Generative AI shock (Table 8) exacerbates these differences. Public firms largely avoided the commoditization of broad adoption; their post-GPT adoption interaction is an insignificant 1.68% penalty. Private firms, however, experienced a clear, significant penalty of 7.07%. Table 9 shows that the adoption penalty for private firms becomes increasingly negative from 2023 to 2025, pointing to a severe reversal in the returns to their earlier upskilling efforts.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.52***	7.01**	7.41***
	(1.66)	(3.15)	(1.67)
AI Specialist	-21.62**	-29.21*	-19.48**
	(8.77)	(16.80)	(8.71)
AI Adopter × GPT	-5.35***	-1.68	-7.07***
	(0.78)	(1.35)	(0.79)
AI Specialist × GPT	15.60***	19.40***	11.92***
	(3.50)	(6.13)	(3.36)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 8: **Impact of ChatGPT on The Relationship Between AI Integration and Productivity by Listing Status**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across publicly listed and private firms. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 7. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills (AI Adoption × GPT) and the complementarity of deep expertise (AI Specialist × GPT). All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 7). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.73***	7.17**	7.51***
	(1.66)	(3.17)	(1.67)
AI Specialist	-22.94***	-30.12*	-20.03**
	(8.82)	(16.96)	(8.77)
AI Adopter × 2023	-4.82***	-1.45	-5.76***
	(0.72)	(1.32)	(0.76)
AI Adopter × 2024	-6.27***	-1.77	-12.10***
	(1.12)	(1.62)	(1.34)
AI Adopter × 2025	-34.15***	-18.27*	-293.17***
	(11.52)	(10.15)	(94.11)
AI Specialist × 2023	10.98***	16.52***	7.92**
	(3.20)	(5.83)	(3.31)
AI Specialist × 2024	24.44***	20.77***	30.20***
	(5.41)	(7.33)	(7.28)
AI Specialist × 2025	154.55***	138.90***	-2630.18*
	(41.62)	(39.34)	(1373.08)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 9: **Temporal Dynamics of the Generative AI Shock by Listing Status**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns for public and private firms. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 7. The interaction terms evaluate the year-by-year acceleration of the commoditization effect for broad skills and the complementarity effect for specialized expertise following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 7). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

At the same time, public firms quickly pivoted to capture the demand for deeper structural integration. The post-GPT interaction for AI specialists in public firms is a robust, significant premium of 19.40%, allowing them to benefit from the new complementarity premium. Private firms also captured a specialist premium, though it was a smaller but still significant 11.92%. Public capital markets seem to provide the necessary financing and governance structures for firms to navigate the commoditization of general AI skills and reallocate resources toward specialized architectural integration.

5 Robustness of Results

To ensure the validity and stability of our primary findings, this section presents two key robustness checks. First, we re-estimate our models on a restricted subsample of highly AI-exposed industries to verify that our results are not driven by broad compositional differences between knowledge-intensive and non-tech sectors. Second, we test the sensitivity of our results to alternative variable constructions by utilizing raw headcounts of AI talent rather than workforce proportions. These checks confirm that our core economic narrative remains robust to both sample restrictions and measurement choices.

5.1 High AI-Exposed Industries

To ensure that the global macroeconomic trends documented above are not driven merely by compositional differences between structurally non-exposed sectors (e.g., agriculture, hospitality) and knowledge-intensive sectors, we re-estimate our models on a highly restricted sub-sample. We filter the Orbis data to strictly retain firms operating within "AI-exposed" or knowledge-intensive industries. Specifically, we retain 2-digit NAICS codes 51 (Information/Tech), 52 (Finance and Insurance), 54 (Professional, Scientific, and Technical Services), 55 (Management of Companies), and 31-33 (Manufacturing). The full set of results for this restricted sample is reported in the Appendix (Tables 10 through 18).

The results from this restricted sample closely mirror the primary findings from Subsections 4.1 through 4.3, confirming that our core economic narrative is driven by fundamental organizational mechanisms rather than sectoral composition. Across the income group dimensions, the baseline return to broad AI adoption remains a positive and significant 6.58% globally. However, the Generative AI shock still exerts a significant commoditization effect

even within these tech-heavy sectors. The post-GPT penalty for broad adoption is a significant 4.53% globally, and this penalty remains particularly severe in Upper-Middle-Income Countries, which face a 27.59% reduction in productivity per percentage point increase in upskilling. Meanwhile, the complementarity premium for AI specialists post-GPT remains strong, yielding a significant 14.15% premium globally and a 15.70% premium in High-Income Countries.

Similarly, the scale advantages observed in the main sample persist within highly exposed industries. Small firms continue to face a significant baseline penalty for broad adoption of 4.70%, which worsens by an additional 10.19% after the GPT shock. Large and very large firms remain relatively insulated from this commoditization. Furthermore, the reward for structural reorganization is still concentrated among the largest firms, with very large enterprises capturing a significant 50.18% premium for AI specialists following the generative shock.

Finally, the corporate governance dynamics remain broadly consistent with the main sample. At baseline, private firms in exposed industries realized a small but significant 3.81% premium for broad AI adoption, whereas public firms showed no significant baseline adoption premium. However, the generative shock reversed these fortunes. Public firms successfully navigated the shock, remaining completely insulated from the commoditization penalty on broad adoption (an insignificant 0.47% penalty) while capturing a significant 16.99% premium for AI specialists. Private firms, conversely, faced a severe and significant 7.01% penalty for broad adoption post-GPT, entirely erasing their baseline advantage. While private firms also captured a specialist premium (10.72%), it was notably smaller than that of their public counterparts. Ultimately, these restricted-sample estimates prove that the divergence between specialized structural integration and superficial technological diffusion is a robust organizational phenomenon, holding true even among firms positioned closest to the technological frontier.

5.2 Raw Headcounts

As a final robustness check, we ensure that our findings are not merely artifacts of the probability-of-representation sampling weights used to correct for online profile bias. To this end, we re-estimate our full suite of models utilizing raw, unweighted headcounts rather

than weighted workforce proportions. Tables 19 through 27 systematically replicate the baseline, Generative AI interaction, and temporal decomposition models under this alternative aggregation specification.

The results utilizing raw headcounts are highly consistent with our primary findings in Subsections 4.1 through 4.3. Across the global sample, the baseline return to broad AI adoption remains significantly positive, yielding a 6.22% increase in productivity globally. Following the Generative AI shock, this broad adoption is commoditized, imposing a significant 4.92% penalty globally and a severe 26.37% penalty in Upper-Middle-Income Countries. Conversely, the complementarity premium for AI specialists emerges robustly post-GPT, generating a significant 13.66% premium globally that is largely driven by High-Income Countries (15.12%).

The scale and governance dynamics also hold firm under this alternative measurement. The temporal decomposition confirms that small firms face an accelerating penalty for broad adoption that deepens further from 2024 to 2025. Conversely, very large firms remain insulated from this adoption penalty and instead capture a massive, expanding premium for dedicated specialists over the same period. As noted previously, the extreme point estimates for 2025 reflect the incomplete financial data cutoff for that year, though the directional intensification remains entirely consistent. Furthermore, public firms effectively avoid the post-GPT commoditization penalty (facing an insignificant 0.85% penalty) and secure a robust 16.59% premium for specialists. Private firms, however, suffer a significant 7.14% penalty on their broad adoption investments.

Ultimately, this alternative specification proves that the aforementioned dichotomies are a fundamental characteristic of the modern firm's production function, entirely robust to the underlying method of headcount aggregation.

6 Conclusion

This paper provides a comprehensive, firm-level examination of Artificial Intelligence as a transformative—yet highly specialized—form of intangible capital. By linking globally representative worker-level microdata with firm financial performance, we document a stark dichotomy in the organizational returns to AI integration. We find that broad technological diffusion—simply upskilling a traditional workforce with AI tools—yields minimal produc-

tivity premiums and, in emerging markets or highly exposed industries, acts as a costly distraction. Conversely, the true economic value of AI is realized almost exclusively along the intensive margin of structural reorganization, where firms explicitly dedicate headcount and redraw boundaries to formalize specialized AI talent.

Furthermore, we demonstrate that the advent of Generative AI acted as a profound structural shock that accelerated these diverging outcomes. Rather than replacing high-end technical labor, highly accessible models like ChatGPT effectively commoditized basic AI literacy while exponentially increasing the premium for deep, specialized architectural integration. This complementarity effect grew significantly stronger over the years following the shock, punishing superficial adoption while generating massive competitive advantages for specialized firms.

Finally, our cross-country decomposition reveals a widening global technological divide. While High-Income Countries possess the complementary infrastructure to absorb and successfully reward both broad and deep AI capital, Lower-Middle-Income Countries face systemic frictions that neutralize productivity gains. Upper-Middle-Income Countries face the most volatile transition, experiencing immense rewards for specialization but devastating penalties for superficial adoption. Ultimately, our findings suggest that navigating the AI intangible shock requires far more than adopting off-the-shelf software; it demands deliberate organizational redesign, structural commitment, and the accumulation of deeply specialized human capital.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, *40*(121), 5–42. https://doi.org/10.3386/w32487

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, *128*(6), 2188–2244. https://doi.org/10.1086/705716

Agrawal, A., Gans, J., & Goldfarb, A. (2019). *The economics of artificial intelligence: An agenda*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.001.0001

Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2023). *Artificial intelligence and firms' systematic risk* [SSRN Working Paper]. https://doi.org/10.2139/ssrn.4868770

Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, *151*, 103745. https://doi.org/10.1016/j.jfineco.2023.103745

Barrios, J. M., Campbell, J. L., Johnson, R. G., & Liu, C. (2025). *Signals or smoke? the determinants and informativeness of corporate artificial intelligence (ai) disclosures* [SSRN Working Paper, last revised May 2025]. https://doi.org/10.2139/ssrn.5133107

Bloom, N., Sadun, R., & Van Reenen, J. (2012). Americans do it better: Us multinationals and the productivity miracle. *American Economic Review*, *102*(1), 167–201. https://doi.org/10.1257/aer.102.1.167

Brynjolfsson, E. (2022). The Turing trap: The promise and peril of human-like artificial intelligence. *Daedalus*, *151*(2), 272–287. https://doi.org/10.1162/daed_a_01915

Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, *13*(1), 333–372. https://doi.org/10.1257/mac.20180386

Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative AI and firm values* (Working Paper No. 31222). National Bureau of Economic Research. https://doi.org/10.3386/w31222

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *Gpts are gpts: An early look at the labor market impact potential of large language models* (Working Paper). OpenAI. https://doi.org/10.48550/arXiv.2303.10130

Hampole, M., Papanikolaou, D., Schmidt, L. D. W., & Seegmiller, B. (2025). *Artificial intelligence and the labor market* (Working Paper No. 33509). National Bureau of Economic Research. https://doi.org/10.3386/w33509

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, *381*(6654), 187–192. https://doi.org/10.1126/science.adh2586

A AI Keyword Categorization

To evaluate the technical depth of AI human capital, we classify raw skill keywords extracted from professional profiles into three distinct categories:

- **Basic AI Keywords:** "AI", "Artificial Intelligence", "ML", "Machine Learning", "GenAI", "Generative AI", "Generative Artificial Intelligence", "Prompt Engineering", "Prompt Engineer", "GPT".
- **Intermediate AI Keywords:** "Deep Learning", "DL", "Neural Network", "Neural Networks", "Neural Net", "NN", "NLP", "Natural Language Processing", "Computer Vision", "AutoML", "Automated Machine Learning", "Gradient Boost", "Gradient Boosting", "XGBoost", "LightGBM", "LLM", "Large Language Model", "Generative Pretrained Transformer", "Generative Pre-Trained Transformer", "Generative Pre Trained Transformer".
- **Advanced AI Keywords:** "CNN", "Convolutional Neural Network", "Convolutional Neural Net", "ConvNet", "RNN", "Recurrent Neural Network", "Recurrent Neural Net", "LSTM", "Long Short-Term Memory", "GAN", "Generative Adversarial Network", "Generative Adversarial Net", "RL", "Reinforcement Learning", "Reinforcement Learn", "TensorFlow", "PyTorch", "Keras", "Hugging Face", "MLflow", "mlflow".

B Matching Revelio to Orbis

We link Revelio companies to Bureau van Dijk's Orbis database using a three-part procedure: (i) text normalization, (ii) name-based candidate retrieval and scoring, and (iii) hierarchical matching on identifiers, name similarity scores, and ancillary attributes. Throughout the process, matching is conducted within country to reduce false positives. The procedure is designed to balance recall (finding true matches even when names differ slightly across sources) with precision (avoiding spurious links). We restrict Orbis matches to entities with valid financial data in at least two years since 2005 to ensure that matched firms are economically meaningful.

B.1 Part 1: Text Normalization

Before matching, we apply a uniform text-cleaning pipeline to company names in both datasets. The procedure has two stages. First, we perform general text normalization: converting all characters to lowercase, decomposing Unicode accented characters to their ASCII base forms (e.g., "é" to "e"), removing diacritical marks, stripping HTML entities, emojis, currency symbols, punctuation, and control characters, and replacing ampersands with "and." Second, we standardize legal entity suffixes to a canonical form. Single-word abbreviations are expanded (e.g., "Co" to "Company," "Corp" to "Corporation," "Ltd" to "Limited," "Inc" to "Incorporated"), while multi-word legal forms are collapsed to their acronym (e.g., "Limited Liability Company" to "LLC," "Public Limited Company" to "PLC"). Leading and trailing instances of "The" are removed. This normalization ensures that superficial formatting differences across data providers do not prevent otherwise identical names from matching.

B.2 Part 2: Name Scoring

We score the similarity between every Revelio company name and Orbis company name within the same country using a GPU-accelerated cascade pipeline. The pipeline consists of three stages, each progressively more precise but also more computationally expensive.

Stage 1: TF-IDF candidate retrieval. We construct character-level n-gram TF-IDF vectors (using trigrams through 5-grams) for all Orbis names within a country. For each Revelio name, we compute cosine similarity against the full Orbis corpus via sparse matrix multiplication on the GPU and retain the top 50 candidates. This step serves as a high-recall filter, efficiently narrowing millions of potential pairs to a tractable candidate set.

Stage 2: Token-set-ratio scoring. Each candidate pair from Stage 1 is scored using the RapidFuzz token-set-ratio metric, a variant of the Levenshtein-based fuzzy matching score that is invariant to word order and robust to partial name overlap. This score provides a complementary similarity signal to the TF-IDF cosine measure.

Stage 3: Phonetic matching. In parallel, we generate Metaphone phonetic encodings for all names in both datasets and merge on exact phonetic key matches. This channel captures sound-alike names that character-level methods may miss (e.g., names transliterated differently across sources).

Stage 4: Cross-encoder reranking. All unique candidate pairs produced by Stages 1–3 are reranked using a pretrained cross-encoder model (MS MARCO MiniLM). The cross-encoder processes each name pair jointly and produces a semantic similarity score, which we normalize to the unit interval via the sigmoid function. This score, which we label *sim*, serves as the primary name-match quality measure used in subsequent steps.

B.3 Part 3: Hierarchical Match Protocol

We assign Revelio–Orbis links through a sequential, priority-based protocol of ten steps. Each Revelio firm is matched at most once: once a firm is linked at a given step, it is removed from the candidate pool for all subsequent steps. Earlier steps impose stricter requirements but yield higher-confidence matches; later steps relax criteria progressively.

Steps 1–3 match on exact shared identifiers. Step 1 merges on Legal Entity Identifier (LEI). Step 2 merges on International Securities Identification Number (ISIN). Step 3 merges on stock ticker jointly with country code. These deterministic links require no name similarity threshold.

Steps 4–8 combine name similarity scores with ancillary attributes. Steps 4 and 5 require a high name similarity score ($sim \geq 0.99$) together with a matching website URL or postal code, respectively. Steps 6 and 7 apply the same attribute checks (URL and postal code) but at a moderately relaxed name threshold ($sim \geq 0.99$ or a RapidFuzz score ≥ 99). Step 8 requires a high name score and a matching phone number. When matching on URLs, we exclude generic social-media domains (e.g., Facebook, LinkedIn) that would generate spurious links.

Steps 9–10 rely on combinations of name similarity and geographic information. Step 9 requires a high name score ($sim \geq 0.99$) together with matching city, country, and a two-digit postal-code prefix. Step 10 links firms whose normalized names are exactly identical and unique within both datasets, conditional on sharing the same country and postal-code prefix.

The final output is a many-to-one mapping from Revelio company identifiers to Orbis BvD identifiers, with a *link_source* flag recording the step at which each match was established.

C Figures: AI Adoption

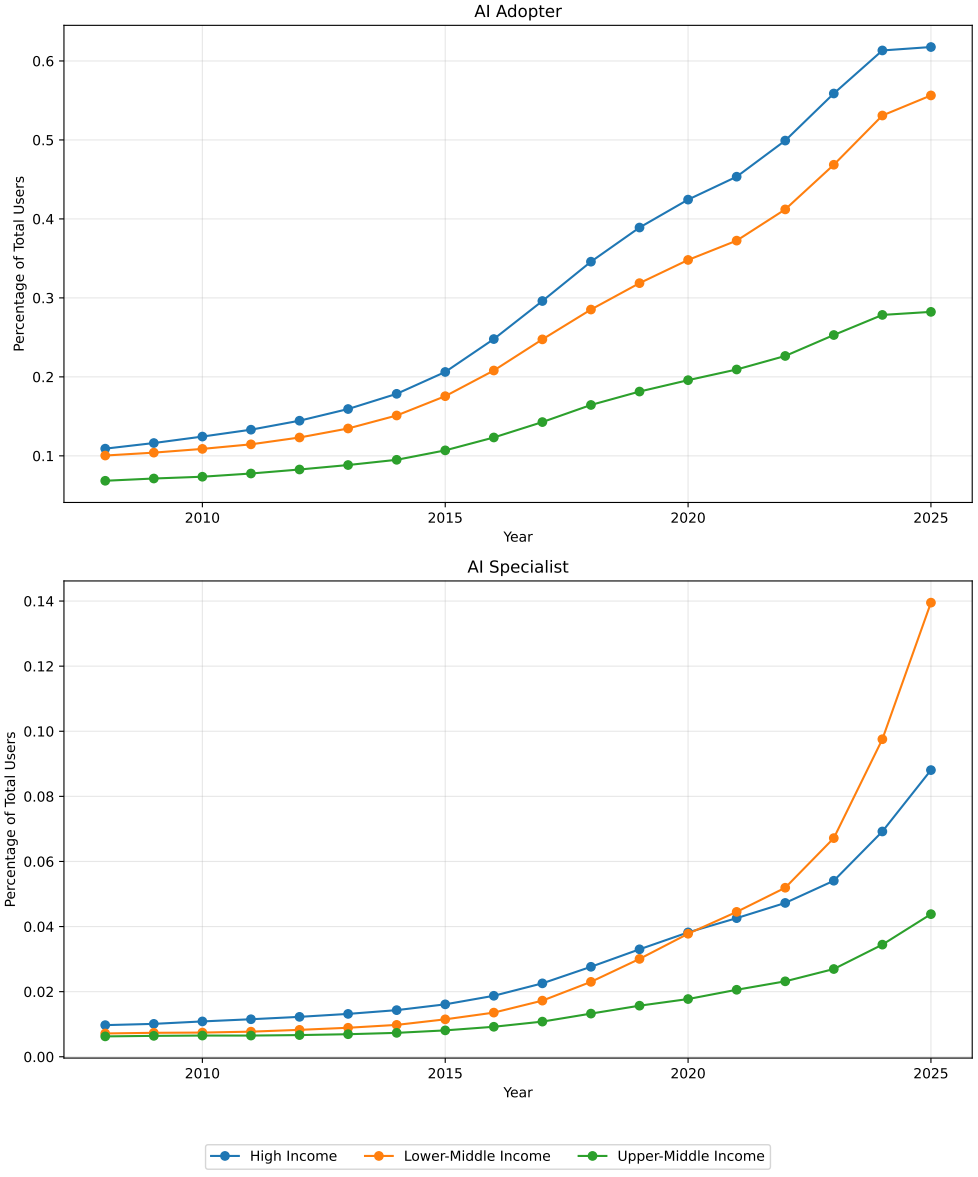


Figure 1: **Annualized AI Density by Income Group (2010-2025)**

This figure illustrates the density of AI-related positions across different World Bank income groups: High Income (HIC), Upper-Middle Income (UMC), Lower-Middle Income (LMC), and Low Income (LIC). The y-axis represents the percentage of total jobs that are tagged as AI Adoption. The data reveals a significant divergence in labor market trends, with High-Income countries experiencing a sharp exponential increase in the share of AI jobs—surpassing 60% by 2025—while Low-Income and Lower-Middle-Income groups show a much slower rate of adoption, highlighting a growing global technological gap.

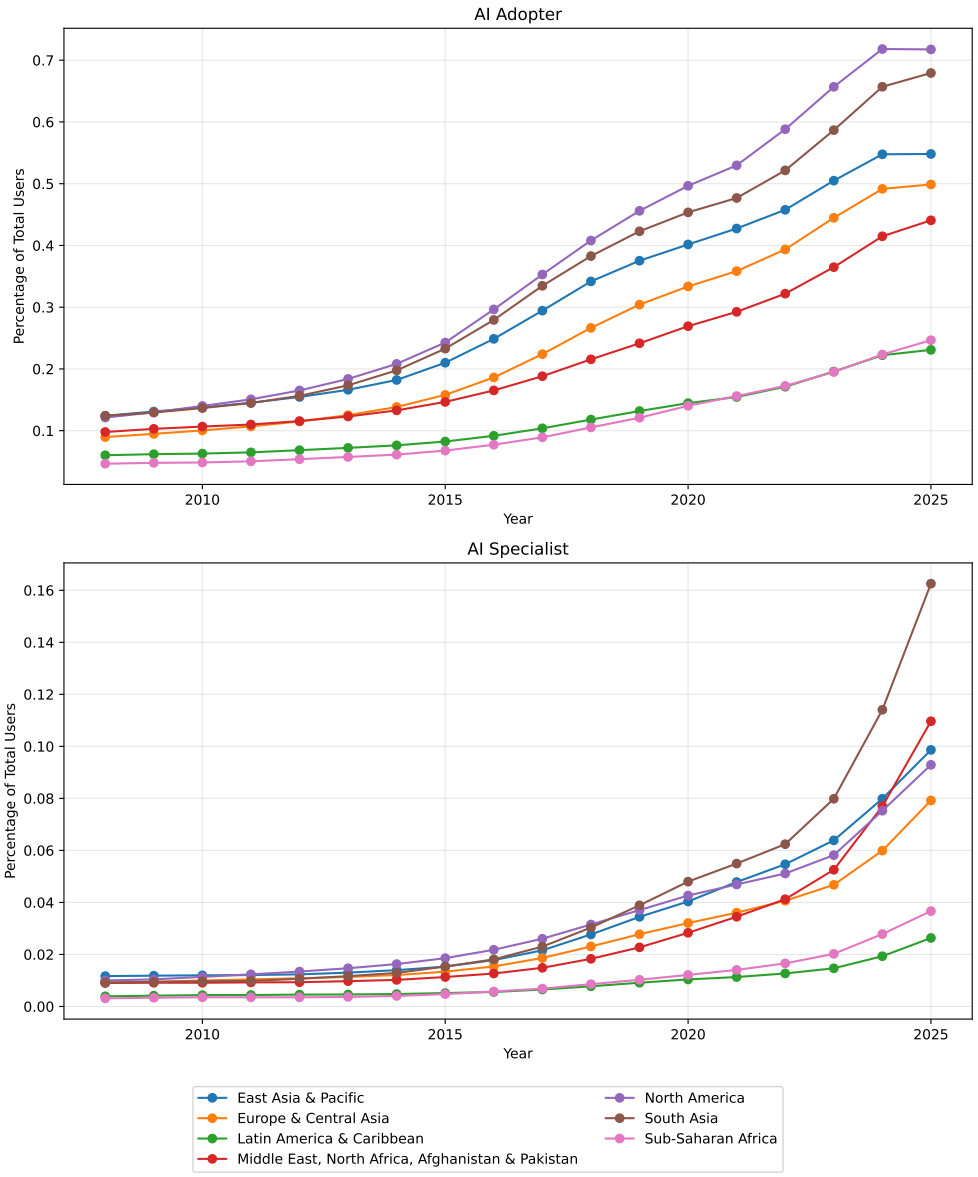


Figure 2: **Annualized AI Density by World Bank Region (2010-2025)**

This figure plots the regional distribution of AI-related employment across seven global regions from 2010 to 2025. The y-axis represents the percentage of total jobs that are tagged as AI Adoption. North America is the clear leader in this transition, with the steepest growth trajectory reaching approximately 80% density by the end of the period. In contrast, regions such as Sub-Saharan Africa and South Asia exhibit the lowest levels of AI integration, with the share of AI-tagged roles remaining significantly lower than in the Global North and East Asia.

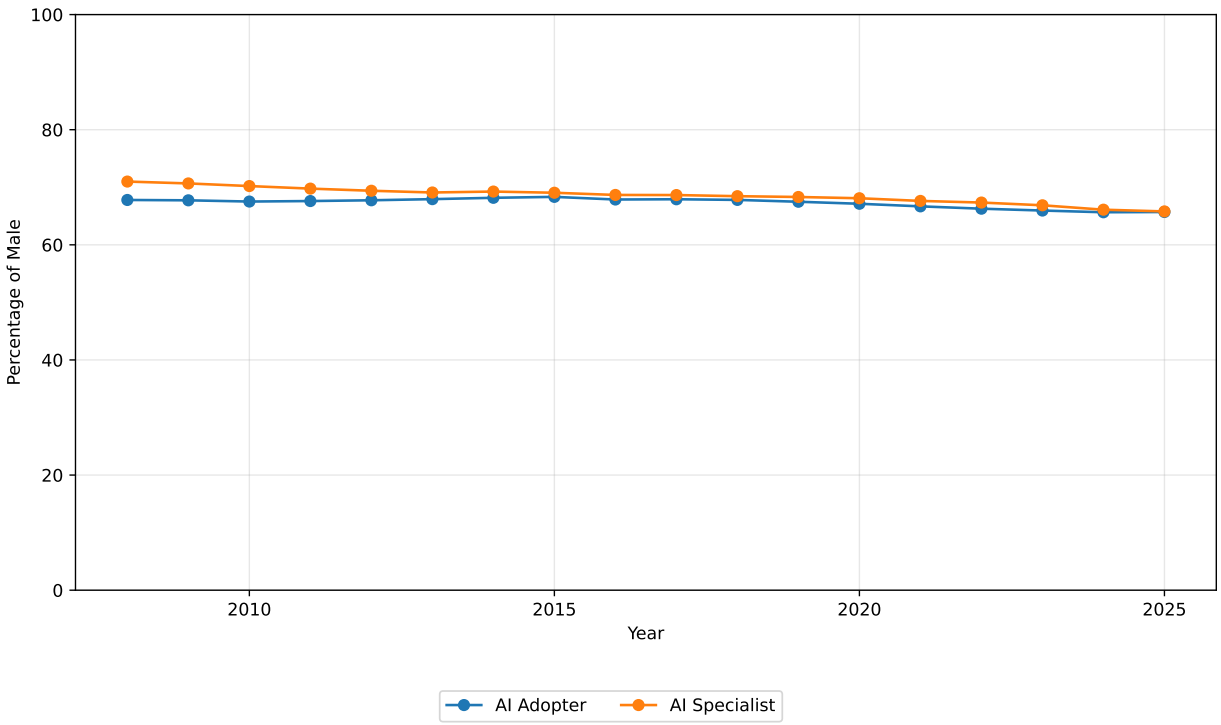


Figure 3: **Predicted Gender Breakdown (2010-2025)**

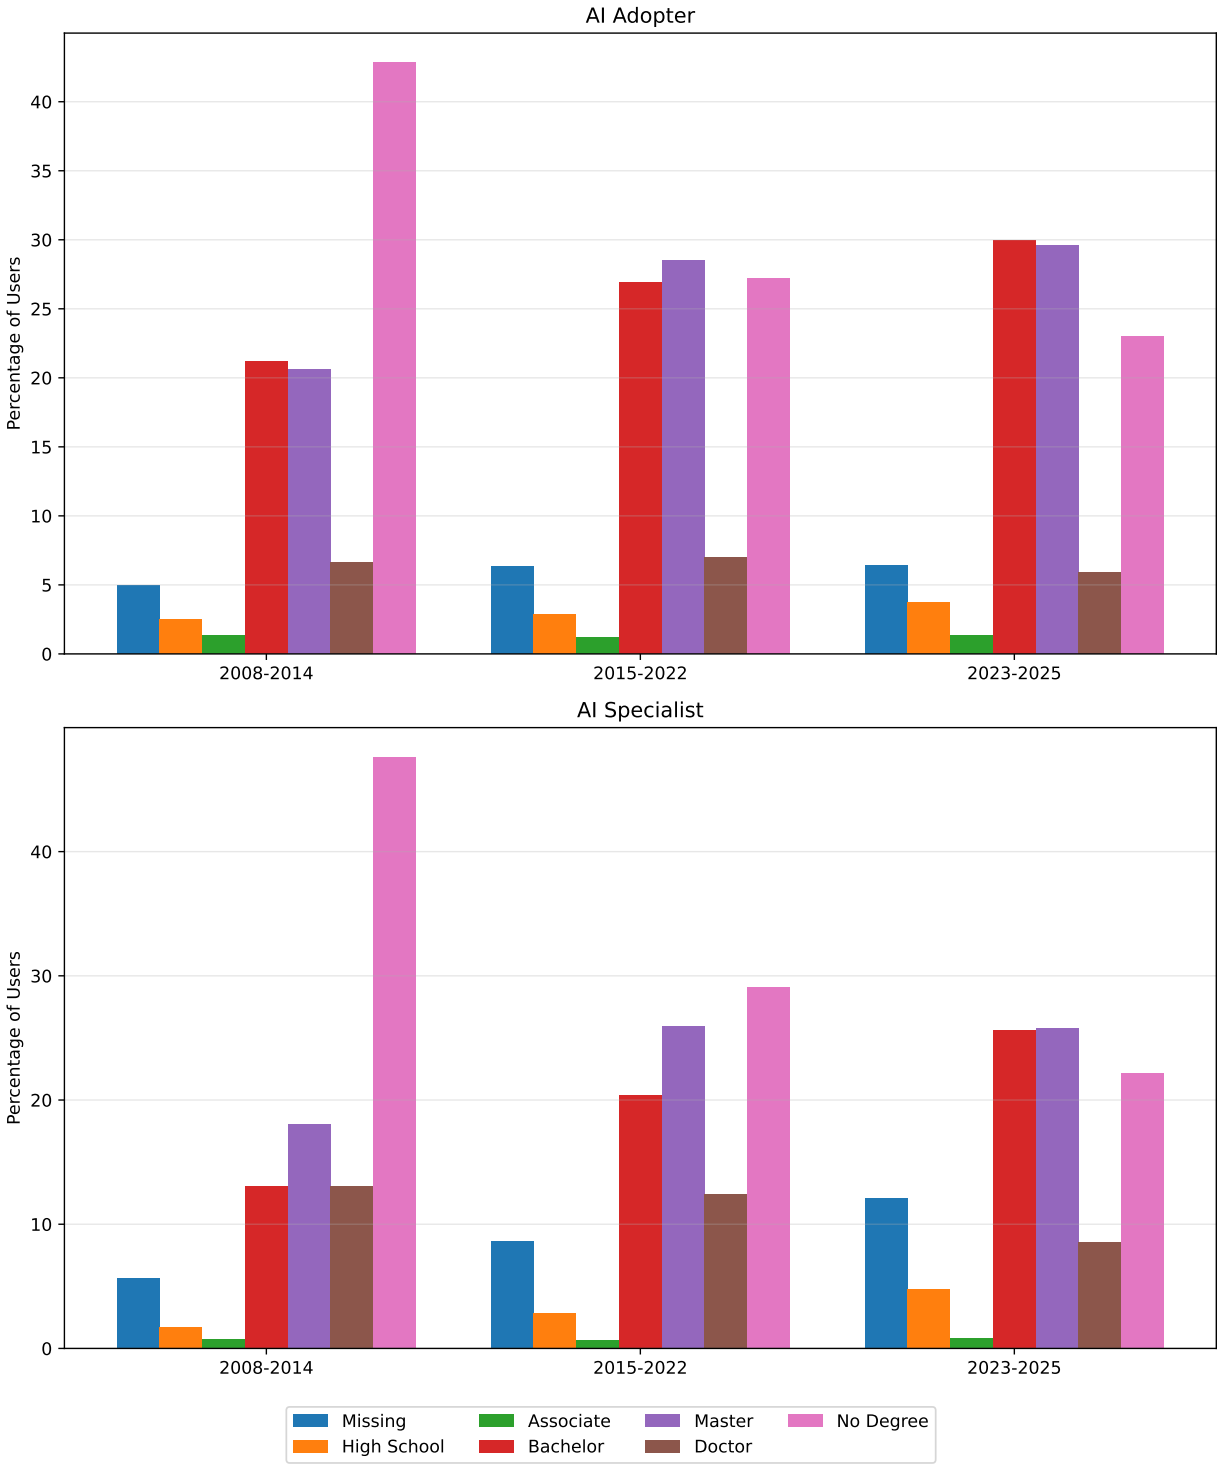


Figure 4: **Distribution of Highest Degree Obtained**

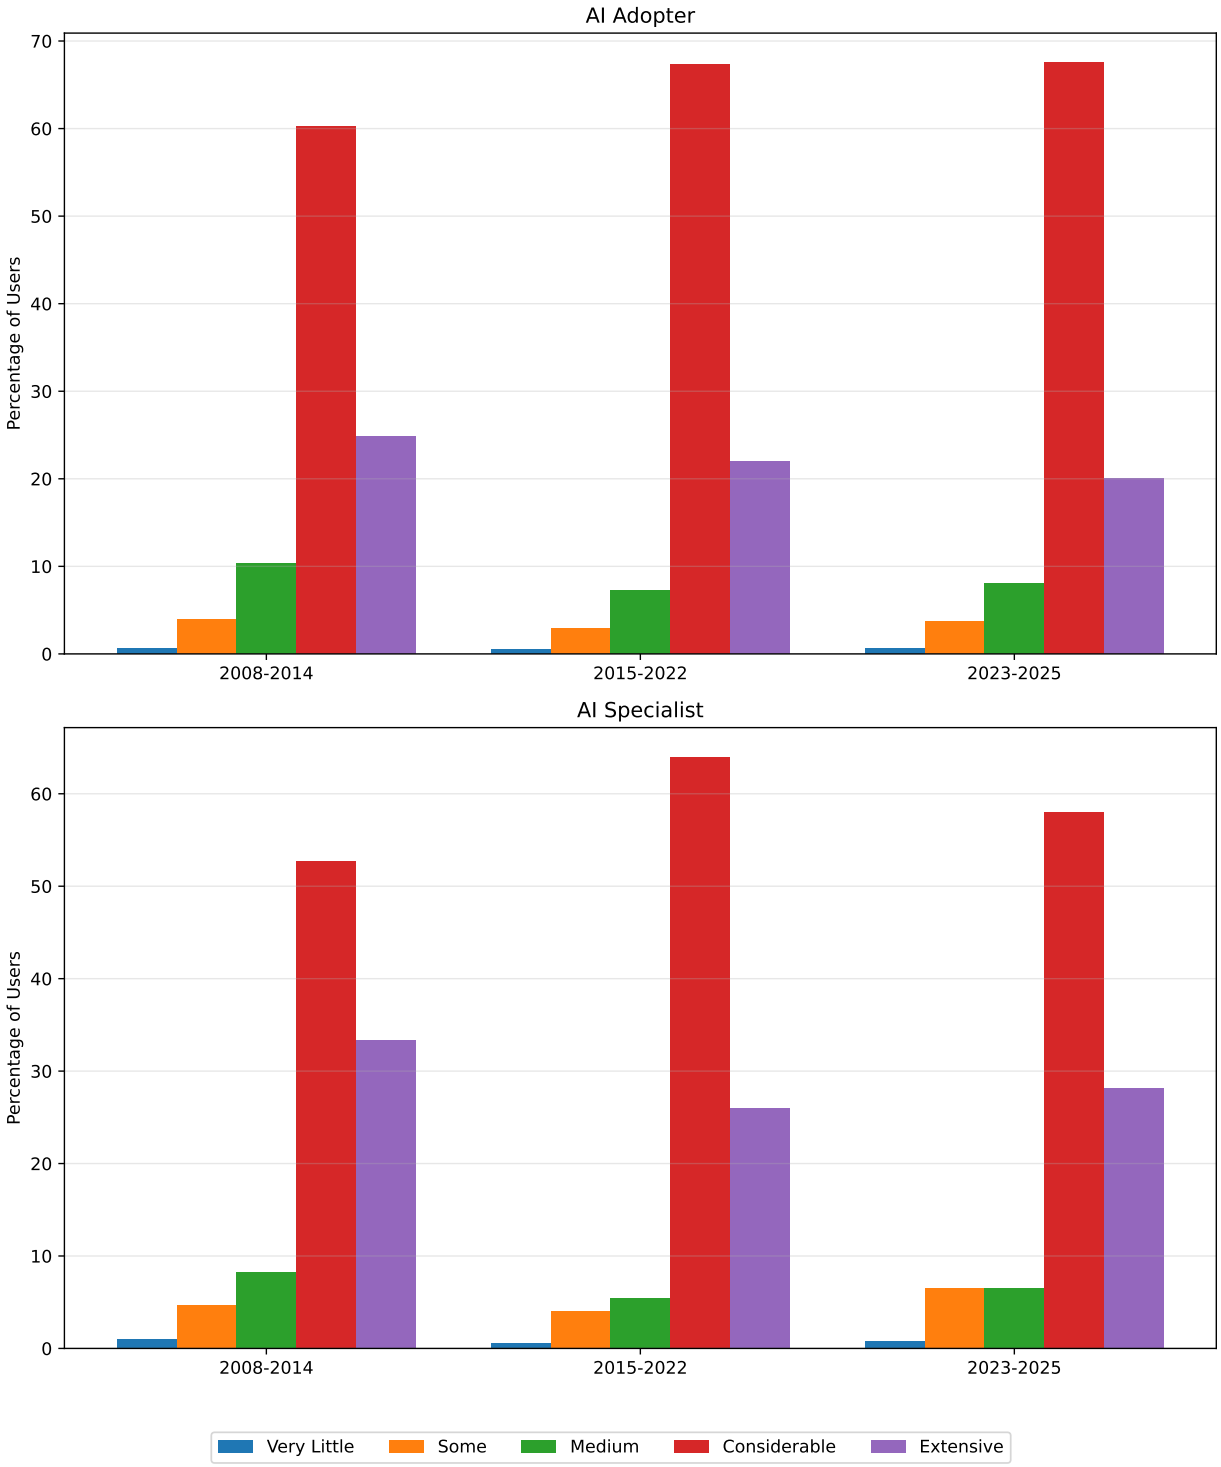


Figure 5: **Distribution of O*Net Job Zone (Preparation Time)**

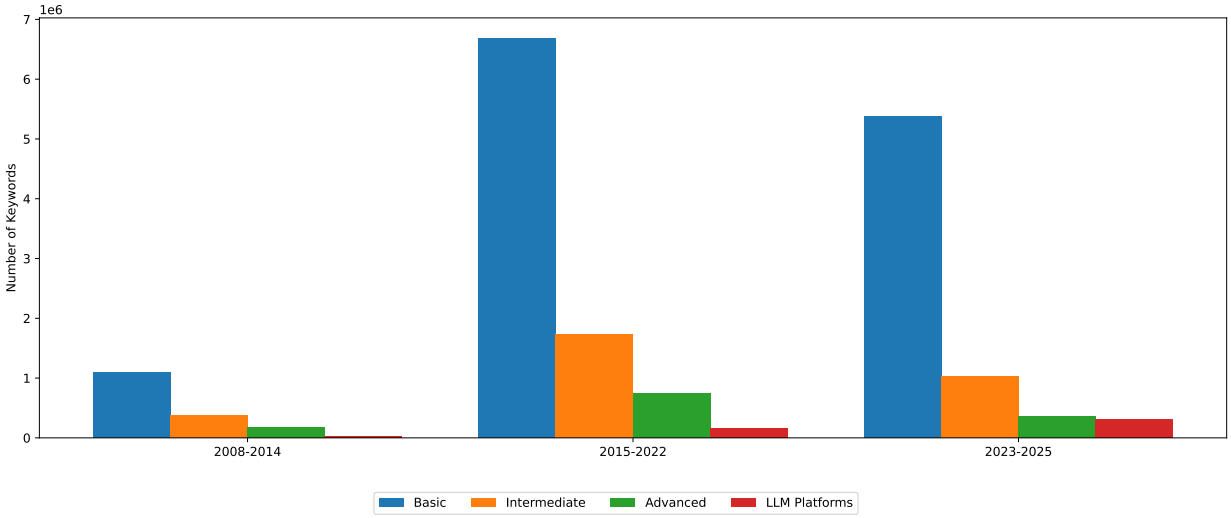


Figure 6: **Distribution of AI-Related Keywords**

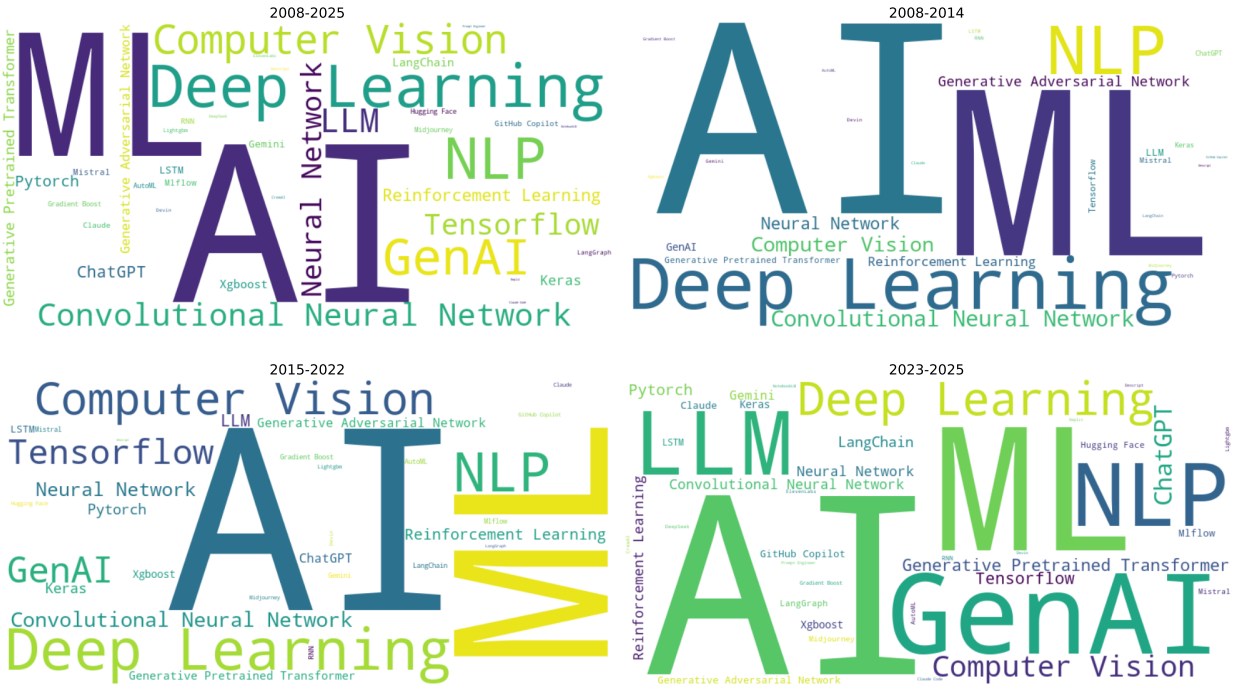


Figure 7: **Word Cloud of AI Keywords**

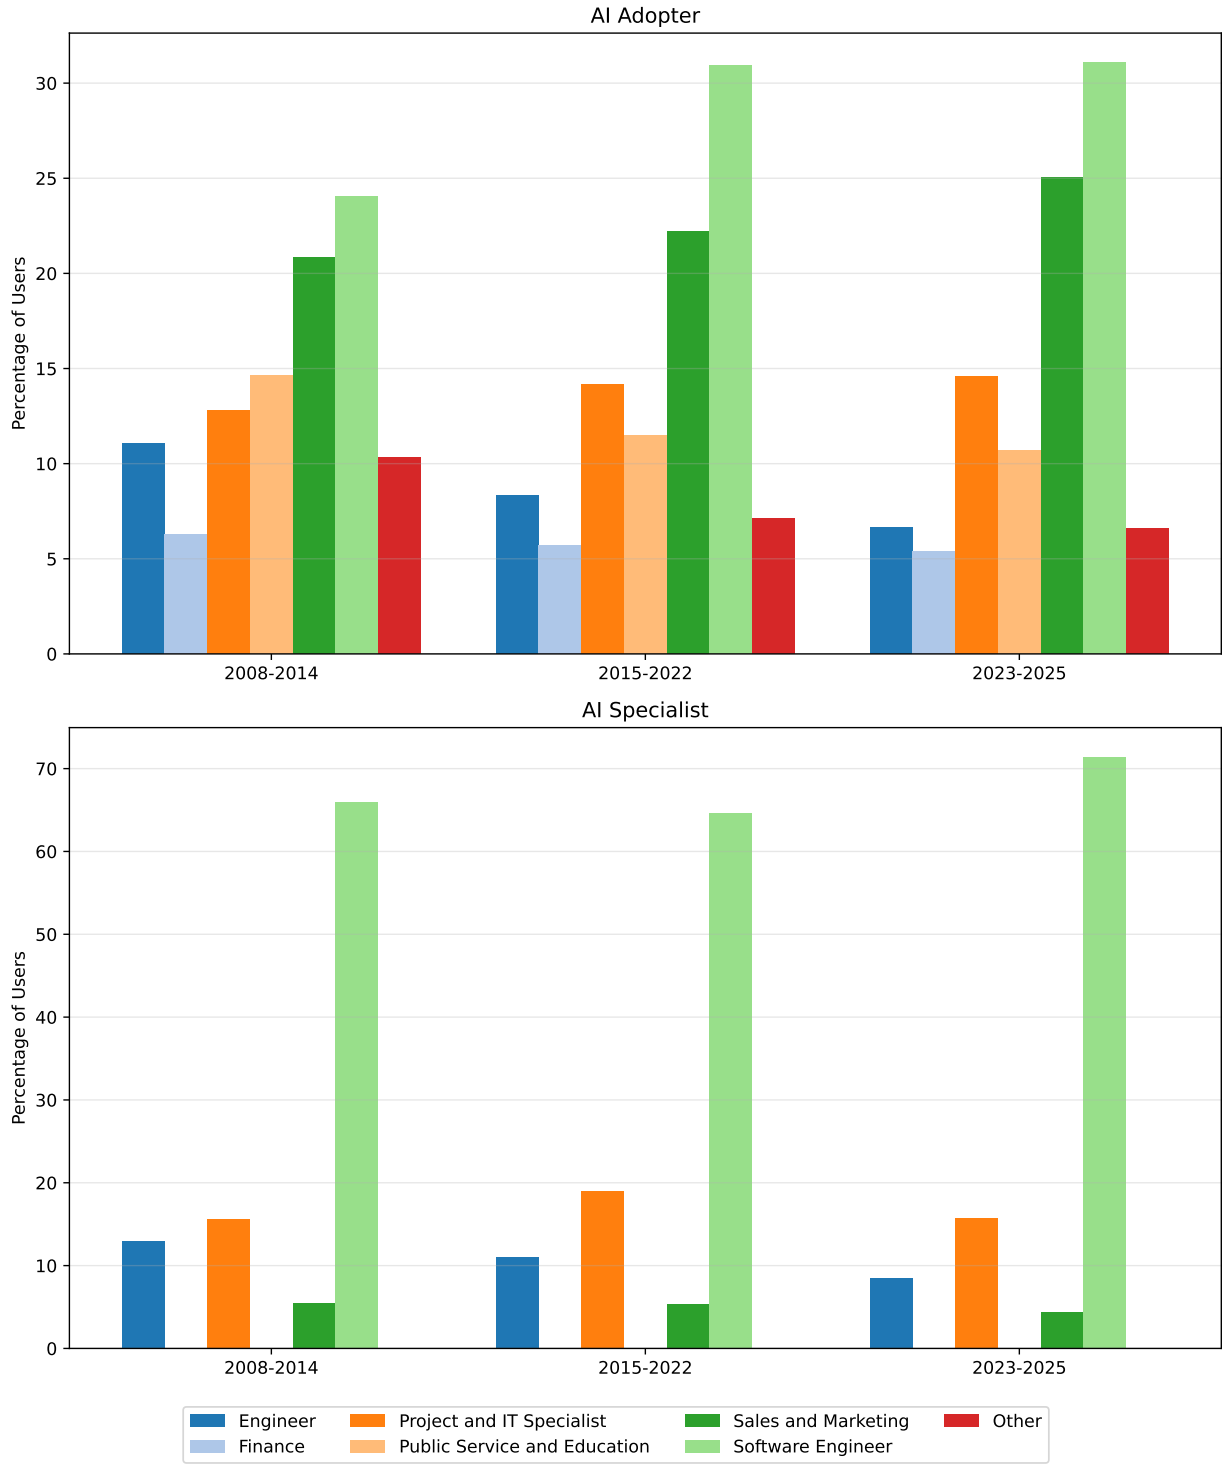


Figure 8: **Distribution of Roles**

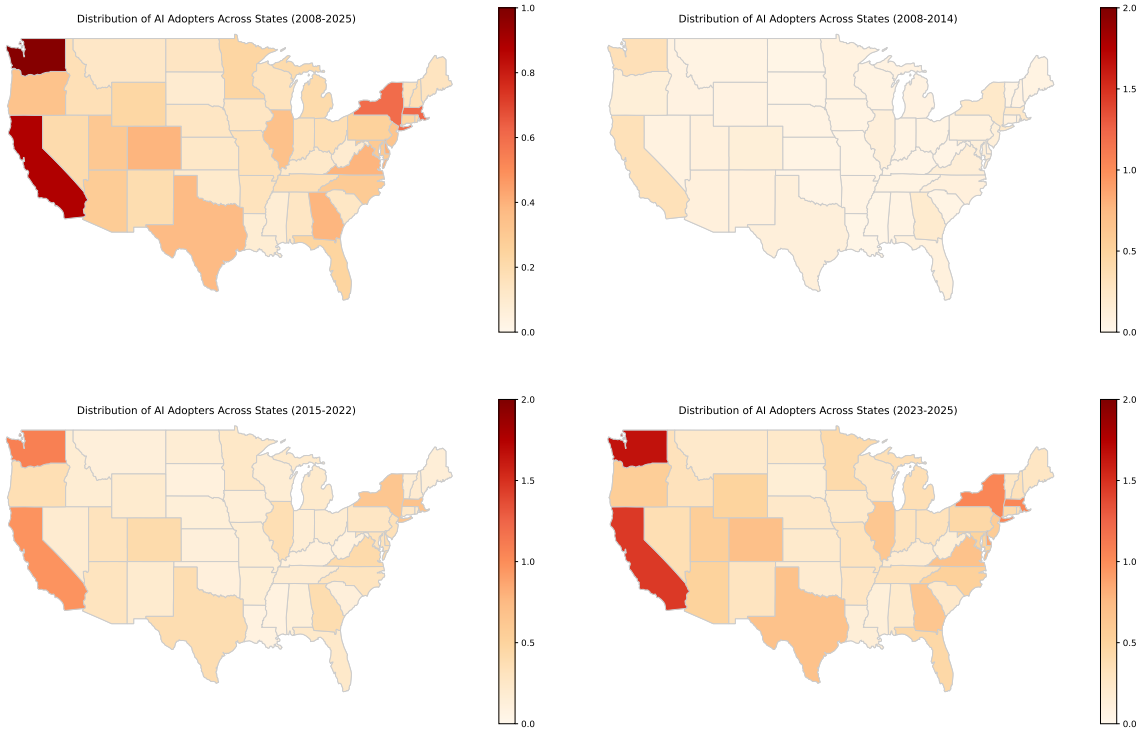


Figure 9: **Distribution of AI Positions Across U.S. States**

This figure presents a geographic heat map of the United States, where the shading intensity represents the percentage of total jobs that are tagged as AI Adoption within each state. The map identifies key clusters of AI activity, with higher concentrations typically observed in technology hubs, indicating an uneven spatial distribution of AI-related labor demand across the country.

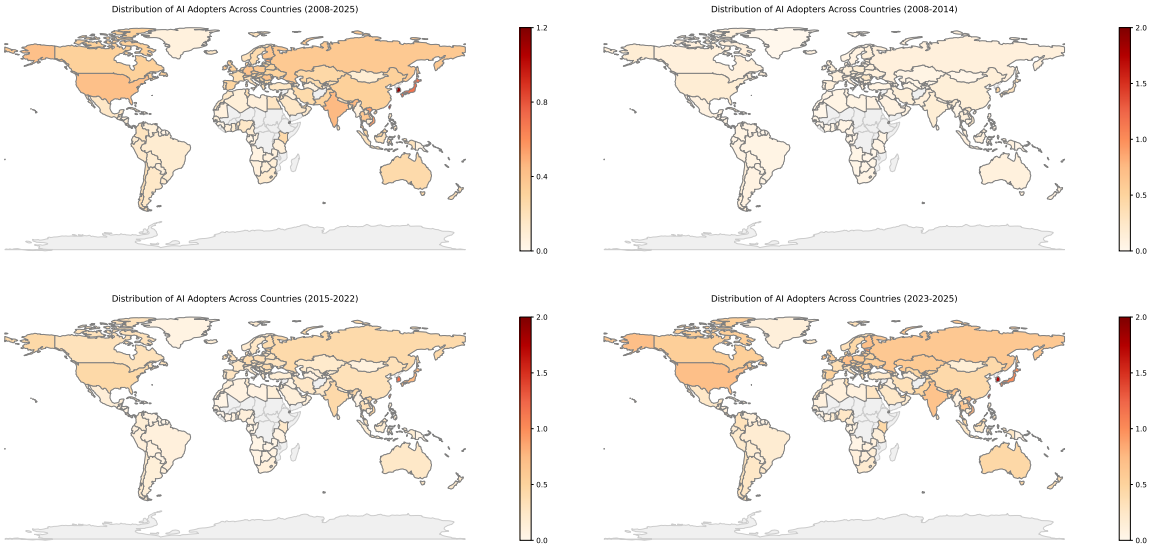


Figure 10: **Global Distribution of AI Positions Across Countries**

This world map illustrates the global landscape of AI employment by country. The color scale indicates the percentage of total jobs that are tagged as AI Adoption. The visualization highlights the dominance of the Global North and specific emerging economies in AI integration, while also showcasing the regions where AI adoption in the formal job market remains in its early stages.

D Highly AI-Exposed Industries

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	6.58***	10.15***	-3.07	47.94***
	(1.85)	(1.87)	(8.87)	(11.35)
AI Specialist	-15.54	-17.70*	86.20	-71.72
	(10.01)	(9.99)	(61.81)	(67.59)
Capital	0.11***	0.08***	0.16***	0.23***
	(0.00)	(0.00)	(0.01)	(0.02)
Labor	-0.30***	-0.29***	-0.26***	-0.39***
	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	327,857	265,973	38,341	23,536
Adjusted R^2	0.926	0.924	0.886	0.925

Table 10: **Baseline Productivity and AI Integration in Highly AI-Exposed Industries**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** is the sampling-weighted fraction of the firm's workforce possessing general AI skills. **AI Specialist** is the sampling-weighted fraction of the workforce holding formal, dedicated AI job titles. The sample is decomposed across World Bank Income Groups: High-Income Countries (HIC), Upper-Middle-Income Countries (UMC), and Lower-Middle-Income Countries (LMC). **Capital** is defined as the logarithm of the ratio of total fixed assets to the total number of employees. **Labor** is defined as the logarithm of the total number of employees. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls (firm size, proxied by total assets; firm age; and overall tech intensity), as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.56***	12.90***	21.89**	58.32***
	(2.01)	(2.05)	(9.26)	(12.81)
AI Specialist	-26.14**	-30.29***	95.55	-65.92
	(10.55)	(10.56)	(66.16)	(78.66)
AI Adopter × GPT	-4.53***	-3.18***	-27.59***	-7.46*
	(0.93)	(0.96)	(3.59)	(4.35)
AI Specialist × GPT	14.15***	15.70***	13.36	-0.83
	(4.10)	(4.27)	(16.21)	(24.47)
Capital	0.11***	0.08***	0.16***	0.23***
	(0.00)	(0.00)	(0.01)	(0.02)
Labor	-0.30***	-0.29***	-0.25***	-0.39***
	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	327,857	265,973	38,341	23,536
Adjusted R^2	0.926	0.924	0.887	0.925

Table 11: **Impact of ChatGPT on Productivity in Highly AI-Exposed Industries**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills (AI Adoption × GPT) and the complementarity of deep expertise (AI Specialist × GPT). All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 10). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.76***	13.12***	22.55**	59.10***
	(2.01)	(2.05)	(9.28)	(13.35)
AI Specialist	-27.31**	-32.25***	95.63	-64.75
	(10.60)	(10.60)	(67.15)	(83.17)
AI Adopter × 2023	-4.12***	-3.86***	-22.86***	-6.70
	(0.87)	(0.91)	(3.41)	(4.19)
AI Adopter × 2024	-5.03***	-1.72	-36.88***	-8.95*
	(1.32)	(1.39)	(5.10)	(5.40)
AI Adopter × 2025	-41.30***	9.11	-97.57	-43.53***
	(13.08)	(6.86)	(108.28)	(15.18)
AI Specialist × 2023	10.28***	11.37***	5.31	-2.17
	(3.76)	(3.95)	(14.52)	(21.73)
AI Specialist × 2024	20.32***	23.82***	26.70	1.56
	(6.15)	(6.47)	(22.49)	(32.65)
AI Specialist × 2025	167.74***	110.70***	5.77	-155.34**
	(44.86)	(31.24)	(131.80)	(62.71)
Capital	0.11***	0.08***	0.16***	0.23***
	(0.00)	(0.00)	(0.01)	(0.02)
Labor	-0.30***	-0.29***	-0.25***	-0.39***
	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	327,857	265,973	38,341	23,536
Adjusted R^2	0.926	0.924	0.887	0.925

Table 12: **Temporal Dynamics of the Generative AI Shock in Highly AI-Exposed Industries**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. The interaction terms evaluate the year-by-year acceleration of the commoditization effect for broad skills and the complementarity effect for specialized expertise following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 10). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	6.58***	-4.70*	11.36***	14.56***	13.72***
	(1.85)	(2.60)	(3.26)	(5.25)	(5.07)
AI Specialist	-15.54	15.29	-4.41	-78.75**	-59.59*
	(10.01)	(19.10)	(14.04)	(32.42)	(31.50)
Capital	0.11***	0.10***	0.09***	0.13***	0.17***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)
Labor	-0.30***	-0.30***	-0.36***	-0.36***	-0.26***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	327,857	154,096	119,394	33,197	16,158
Adjusted R^2	0.926	0.934	0.925	0.934	0.954

Table 13: **Baseline Productivity and AI Integration by Firm Size in Highly AI-Exposed Industries**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by firm size within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. The sample is decomposed by firm size categories: Small, Medium, Large, and Very Large. **Capital** and **Labor** are defined as in Table 10. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls, as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.56***	3.64	14.03***	16.41***	16.07***
	(2.01)	(2.73)	(3.65)	(5.82)	(5.67)
AI Specialist	-26.14**	8.11	-4.95	-87.98**	-124.41***
	(10.55)	(19.72)	(14.93)	(35.23)	(34.06)
AI Adopter × GPT	-4.53***	-10.19***	-3.29*	-1.74	-0.74
	(0.93)	(1.10)	(1.82)	(2.59)	(2.33)
AI Specialist × GPT	14.15***	15.48***	2.57	9.20	50.18***
	(4.10)	(5.56)	(6.37)	(10.64)	(15.97)
Capital	0.11***	0.10***	0.09***	0.13***	0.17***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)
Labor	-0.30***	-0.30***	-0.36***	-0.36***	-0.26***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	327,857	154,096	119,394	33,197	16,158
Adjusted R^2	0.926	0.934	0.925	0.934	0.954

Table 14: **Impact of ChatGPT on Productivity by Firm Size in Highly AI-Exposed Industries**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across different firm size categories within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills and the complementarity of deep expertise. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.76***	3.96	14.11***	16.65***	16.30***
	(2.01)	(2.73)	(3.65)	(5.83)	(5.67)
AI Specialist	-27.31**	7.20	-4.95	-92.04**	-127.45***
	(10.60)	(19.61)	(15.00)	(35.82)	(33.52)
AI Adopter × 2023	-4.12***	-7.82***	-3.65**	-2.49	-2.51
	(0.87)	(1.09)	(1.69)	(2.56)	(2.22)
AI Adopter × 2024	-5.03***	-14.88***	-2.15	-0.15	1.78
	(1.32)	(1.53)	(2.66)	(3.61)	(2.78)
AI Adopter × 2025	-41.30***	-138.97***	-129.12***	-27.69	1.02
	(13.08)	(42.31)	(35.07)	(32.75)	(7.13)
AI Specialist × 2023	10.28***	9.84*	3.64	6.36	43.31***
	(3.76)	(5.38)	(5.85)	(10.71)	(14.25)
AI Specialist × 2024	20.32***	25.71***	-0.56	13.10	51.35***
	(6.15)	(9.70)	(10.37)	(13.85)	(18.93)
AI Specialist × 2025	167.74***	22.98	718.49***	137.89*	129.50***
	(44.86)	(164.37)	(268.99)	(81.97)	(39.41)
Capital	0.11***	0.10***	0.09***	0.13***	0.17***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)
Labor	-0.30***	-0.30***	-0.36***	-0.36***	-0.26***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	327,857	154,096	119,394	33,197	16,158
Adjusted R^2	0.926	0.934	0.925	0.934	0.954

Table 15: **Temporal Dynamics of the Generative AI Shock by Firm Size in Highly AI-Exposed Industries**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns across firm size classifications within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. The interaction terms evaluate the year-by-year acceleration of the commoditization and complementarity effects. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	6.58***	2.09	3.81**
	(1.85)	(3.24)	(1.83)
AI Specialist	-15.54	-10.98	-17.64*
	(10.01)	(16.79)	(10.52)
Capital	0.11***	0.17***	0.00
	(0.00)	(0.01)	(0.00)
Labor	-0.30***	-0.23***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	327,857	102,805	225,050
Adjusted R^2	0.926	0.898	0.950

Table 16: **Baseline Productivity and AI Integration by Listing Status in Highly AI-Exposed Industries**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by listing status within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. The sample is decomposed into Public and Private firms. **Capital** and **Labor** are defined as in Table 10. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls, as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.56***	2.93	9.07***
	(2.01)	(3.57)	(2.05)
AI Specialist	-26.14**	-31.29*	-23.81**
	(10.55)	(18.22)	(11.06)
AI Adopter × GPT	-4.53***	-0.47	-7.01***
	(0.93)	(1.52)	(0.95)
AI Specialist × GPT	14.15***	16.99**	10.72***
	(4.10)	(6.69)	(4.04)
Capital	0.11***	0.17***	0.00
	(0.00)	(0.01)	(0.00)
Labor	-0.30***	-0.23***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	327,857	102,805	225,050
Adjusted R^2	0.926	0.898	0.950

Table 17: **Impact of ChatGPT on Productivity by Listing Status in Highly AI-Exposed Industries**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across publicly listed and private firms within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills and the complementarity of deep expertise. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.76***	3.09	9.08***
	(2.01)	(3.59)	(2.06)
AI Specialist	-27.31**	-32.06*	-23.93**
	(10.60)	(18.37)	(11.12)
AI Adopter × 2023	-4.12***	-0.42	-5.89***
	(0.87)	(1.49)	(0.93)
AI Adopter × 2024	-5.03***	-0.28	-11.12***
	(1.32)	(1.84)	(1.58)
AI Adopter × 2025	-41.30***	-22.47*	-298.12***
	(13.08)	(11.62)	(55.71)
AI Specialist × 2023	10.28***	14.92**	7.60*
	(3.76)	(6.37)	(4.02)
AI Specialist × 2024	20.32***	17.10**	23.95***
	(6.15)	(7.99)	(8.81)
AI Specialist × 2025	167.74***	144.70***	-985.16***
	(44.86)	(40.59)	(282.95)
Capital	0.11***	0.17***	0.00
	(0.00)	(0.01)	(0.00)
Labor	-0.30***	-0.23***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	327,857	102,805	225,050
Adjusted R^2	0.926	0.898	0.950

Table 18: **Temporal Dynamics of the Generative AI Shock by Listing Status in Highly AI-Exposed Industries**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns for public and private firms within highly AI-exposed industries. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 10. The interaction terms evaluate the year-by-year acceleration of the commoditization and complementarity effects following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Raw Headcounts Robustness Check

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	6.22***	10.02***	3.83	48.29***
	(1.51)	(1.50)	(7.36)	(11.24)
AI Specialist	-12.77	-17.43**	98.85*	-93.92
	(8.11)	(8.05)	(51.89)	(62.80)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.38***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 19: **Baseline Productivity and AI Integration using Raw Headcounts**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** is the raw headcount of the firm's workforce possessing general AI skills. **AI Specialist** is the raw headcount of the workforce holding formal, dedicated AI job titles. The sample is decomposed across World Bank Income Groups: High-Income Countries (HIC), Upper-Middle-Income Countries (UMC), and Lower-Middle-Income Countries (LMC). **Capital** is defined as the logarithm of the ratio of total fixed assets to the total number of employees. **Labor** is defined as the logarithm of the total number of employees. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls (firm size, proxied by total assets; firm age; and overall tech intensity), as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.52***	13.42***	21.91***	54.44***
	(1.64)	(1.66)	(7.67)	(12.69)
AI Specialist	-22.60***	-29.00***	105.93*	-96.08
	(8.58)	(8.55)	(55.48)	(71.88)
AI Adopter × GPT	-4.92***	-3.97***	-26.37***	-4.74
	(0.78)	(0.80)	(3.50)	(4.03)
AI Specialist × GPT	13.66***	15.12***	17.80	4.49
	(3.39)	(3.49)	(15.45)	(19.26)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 20: **Impact of ChatGPT on Productivity using Raw Headcounts**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills (AI Adoption × GPT) and the complementarity of deep expertise (AI Specialist × GPT). All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 19). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity			
	Global	HIC	UMC	LMC
AI Adopter	10.71***	13.61***	22.30***	54.73***
	(1.65)	(1.66)	(7.70)	(13.21)
AI Specialist	-23.75***	-30.79***	105.79*	-97.92
	(8.63)	(8.58)	(55.83)	(75.83)
AI Adopter × 2023	-4.46***	-4.37***	-22.37***	-4.48
	(0.72)	(0.75)	(3.34)	(3.86)
AI Adopter × 2024	-5.67***	-3.11***	-34.36***	-5.20
	(1.14)	(1.19)	(4.86)	(5.06)
AI Adopter × 2025	-34.63***	7.47	-176.86	-48.29***
	(11.52)	(6.16)	(116.83)	(17.29)
AI Specialist × 2023	9.50***	10.49***	8.47	2.33
	(3.10)	(3.21)	(13.80)	(17.25)
AI Specialist × 2024	21.68***	25.29***	35.05	8.79
	(5.29)	(5.53)	(22.26)	(25.32)
AI Specialist × 2025	148.97***	105.42***	53.33	-123.08**
	(40.60)	(30.09)	(133.01)	(58.46)
Capital	0.10***	0.08***	0.17***	0.24***
	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.33***	-0.31***	-0.37***
	(0.01)	(0.01)	(0.01)	(0.02)
Firm FE	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	560,288	471,124	57,623	31,524
Adjusted R^2	0.939	0.943	0.901	0.922

Table 21: **Temporal Dynamics of the Generative AI Shock using Raw Headcounts**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. The interaction terms evaluate the year-by-year acceleration of the commoditization effect for broad skills and the complementarity effect for specialized expertise following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls (including Labor and Capital as defined in Table 19). Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	6.22***	-3.28	9.68***	15.32***	20.14***
	(1.51)	(2.03)	(2.72)	(4.55)	(4.55)
AI Specialist	-12.77	0.01	0.75	-65.21**	-45.15
	(8.11)	(14.01)	(11.57)	(27.24)	(28.47)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.946	0.938	0.943	0.957

Table 22: **Baseline Productivity and AI Integration by Firm Size using Raw Headcounts**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by firm size utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. The sample is decomposed by firm size categories: Small, Medium, Large, and Very Large. **Capital** and **Labor** are defined as in Table 19. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls, as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.52***	4.98**	12.37***	17.71***	23.41***
	(1.64)	(2.15)	(3.06)	(5.05)	(5.16)
AI Specialist	-22.60***	-9.56	2.47	-74.85***	-108.45***
	(8.58)	(14.61)	(12.51)	(29.05)	(31.34)
AI Adopter × GPT	-4.92***	-10.27***	-3.36**	-2.32	-1.82
	(0.78)	(0.92)	(1.57)	(2.25)	(2.04)
AI Specialist × GPT	13.66***	17.38***	0.07	10.25	48.85***
	(3.39)	(4.55)	(5.38)	(9.02)	(14.47)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.946	0.938	0.943	0.958

Table 23: **Impact of ChatGPT on Productivity by Firm Size using Raw Headcounts**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across different firm size categories utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills and the complementarity of deep expertise. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity				
	Global	Small	Medium	Large	Very Large
AI Adopter	10.71***	5.27**	12.45***	17.91***	23.44***
	(1.65)	(2.15)	(3.06)	(5.06)	(5.16)
AI Specialist	-23.75***	-10.72	2.39	-78.35***	-110.87***
	(8.63)	(14.59)	(12.56)	(29.40)	(30.78)
AI Adopter × 2023	-4.46***	-8.17***	-3.49**	-2.83	-3.82*
	(0.72)	(0.90)	(1.44)	(2.20)	(1.95)
AI Adopter × 2024	-5.67***	-14.89***	-2.75	-1.22	0.93
	(1.14)	(1.34)	(2.34)	(3.16)	(2.42)
AI Adopter × 2025	-34.63***	-125.56***	-131.22***	-24.02	0.95
	(11.52)	(34.88)	(35.14)	(29.36)	(5.92)
AI Specialist × 2023	9.50***	12.26***	0.29	6.53	43.26***
	(3.10)	(4.46)	(4.90)	(9.02)	(12.65)
AI Specialist × 2024	21.68***	27.90***	-0.94	16.49	48.90***
	(5.29)	(8.12)	(8.98)	(11.95)	(17.05)
AI Specialist × 2025	148.97***	163.58	705.04**	128.30	129.42***
	(40.60)	(112.04)	(282.96)	(82.69)	(39.16)
Capital	0.10***	0.08***	0.09***	0.14***	0.17***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Labor	-0.33***	-0.35***	-0.39***	-0.37***	-0.31***
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Firm FE	✓	✓	✓	✓	✓
Industry-Year FE	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓
Observations	560,288	278,678	192,150	53,984	27,132
Adjusted R^2	0.939	0.947	0.938	0.943	0.958

Table 24: **Temporal Dynamics of the Generative AI Shock by Firm Size using Raw Headcounts**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns across firm size classifications utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. The interaction terms evaluate the year-by-year acceleration of the commoditization and complementarity effects. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	6.22***	5.04*	2.83*
	(1.51)	(2.81)	(1.46)
AI Specialist	-12.77	-13.10	-12.76
	(8.11)	(15.50)	(7.98)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 25: **Baseline Productivity and AI Integration by Listing Status using Raw Headcounts**

This table presents the baseline fixed-effects estimates evaluating the direct impact of AI integration on firm-level productivity partitioned by listing status utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. The sample is decomposed into Public and Private firms. **Capital** and **Labor** are defined as in Table 19. All specifications include firm fixed effects, industry-by-year fixed effects, and other standard firm-level controls, as indicated by a check mark (✓). These coefficients are not reported to save space. Standard errors are two-way clustered at the firm level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.52***	6.25**	8.28***
	(1.64)	(3.13)	(1.63)
AI Specialist	-22.60***	-32.51*	-20.06**
	(8.58)	(16.84)	(8.49)
AI Adopter × GPT	-4.92***	-0.85	-7.14***
	(0.78)	(1.37)	(0.78)
AI Specialist × GPT	13.66***	16.59***	11.83***
	(3.39)	(5.99)	(3.22)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 26: **Impact of ChatGPT on Productivity by Listing Status using Raw Headcounts**

This table examines how the release of Generative AI (ChatGPT) altered the organizational returns to technology across publicly listed and private firms utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. **GPT** is a binary indicator equal to one for the post-Generative AI boom period ($t \geq 2023$) and zero otherwise. The interaction terms evaluate the commoditization of basic AI skills and the complementarity of deep expertise. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Firm Productivity		
	Global	Public	Private
AI Adopter	10.71***	6.41**	8.34***
	(1.65)	(3.15)	(1.64)
AI Specialist	-23.75***	-33.51**	-20.49**
	(8.63)	(16.98)	(8.54)
AI Adopter × 2023	-4.46***	-0.64	-5.90***
	(0.72)	(1.33)	(0.75)
AI Adopter × 2024	-5.67***	-0.92	-12.08***
	(1.14)	(1.64)	(1.36)
AI Adopter × 2025	-34.63***	-18.09*	-370.87***
	(11.52)	(10.16)	(88.30)
AI Specialist × 2023	9.50***	13.37**	8.22***
	(3.10)	(5.68)	(3.17)
AI Specialist × 2024	21.68***	18.38**	28.45***
	(5.29)	(7.19)	(6.98)
AI Specialist × 2025	148.97***	133.12***	-2100.22**
	(40.60)	(38.12)	(984.10)
Capital	0.10***	0.17***	0.00**
	(0.00)	(0.01)	(0.00)
Labor	-0.33***	-0.26***	-0.57***
	(0.01)	(0.01)	(0.01)
Firm FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
Controls	✓	✓	✓
Observations	560,288	143,822	416,453
Adjusted R^2	0.939	0.911	0.957

Table 27: **Temporal Dynamics of the Generative AI Shock by Listing Status using Raw Headcounts**

This table decomposes the post-Generative AI boom period into distinct annual interaction terms for 2023, 2024, and 2025 to examine the dynamic evolution of technological returns for public and private firms utilizing raw headcounts. The **dependent variable** is the natural logarithm of revenue per employee (in millions of dollars). **AI Adoption** and **AI Specialist** are as defined in Table 19. The interaction terms evaluate the year-by-year acceleration of the commoditization and complementarity effects following the release of ChatGPT. All specifications include firm fixed effects, industry-by-year fixed effects, and the full set of firm-level controls. Standard errors are two-way clustered at the firm and industry level and are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.