

第一页为封面页

参赛学生姓名：阎立谦，邱梓淳，陈胤同

中学：上海星河湾双语学校

省份：上海

国家/地区：中国/南部赛区

指导老师姓名：邓鸣茂 陈晶萍

指导老师单位：上海对外经贸大学

论文题目：Firm-Level Impacts of Artificial Intelligence on Labor Markets: Evidence from Online Job Postings

Firm-Level Impacts of Artificial Intelligence on Labor Markets: Evidence from Online Job Postings

Liqian Yan, Zichun Qiu, Yingtong Chen

Shanghai Starriver Bilingual School

Abstract

This paper explores the influence of Artificial Intelligence (AI) adoption on labor demand creation and displacement in the global economy. We extract firm-level data from 2015 to 2024, which encompasses numerous job categories. Using OLS regression, Instrumental Variable and Difference-in-Differences (DID) method to alleviate endogeneity, the results show that firms adopting AI technologies tend to significantly reduce their reliance on routine labor, while simultaneously reallocating resources toward more complex and creative roles. Furthermore, we find that AI adoption correlates with a significant increase in average offered wages consistent with skill-biased demand and accompanied by the emergence of new jobs that did not appear in our dataset prior to 2018. Our study addresses the absence of firm-level microanalysis and a theoretical model to specifically test our hypotheses regarding AI's impacts on labor: "substitution effect", "complementarity effect", and "wage polarization effect". Finally, we argue that greater attention should be paid to the employment creation effect of artificial intelligence (AI), and guide AI to develop in a positive direction toward creating more new occupations and jobs.

Keywords: Artificial Intelligence (AI); Labor Demand; DID; Job Postings

Contents

1	Introduction	4
2	Literature Review and Contribution	4
2.1	Technological Progress and the Labor Market	4
2.2	The Impact of Artificial Intelligence on Labor Demand	4
2.3	Generative AI and Large Language Models	5
2.4	Heterogeneous Impacts	5
2.5	Our Findings and Contributions	5
3	Theoretical Framework	6
3.1	A Functional Definition of AI	6
3.2	Theoretical Analysis	7
3.2.1	Substitution channel	7
3.2.2	Complementarity channel	7
3.2.3	Job Polarization	7
3.3	A Task-Based Model of AI and Labor	8
3.3.1	The Economic Environment	8
3.3.2	The Firm's Optimization Problem	9
3.3.3	From Model to Measurement: Deriving Testable Hypotheses	10
3.4	Numerical Simulation	11
3.4.1	Calibration	11
3.4.2	Simulation Results	11
3.4.3	Sensitivity Analysis	12
4	Sample Data Selection and Research Design	13
4.1	The Job Postings Universe	13
4.2	Measuring the Independent Variable: A Firm-Level AI Intensity Index	14
4.3	Measuring Dependent Variables: The Task Content of Labor Demand	16
4.4	Statistical Findings	16
4.5	The Correlation Between AI and Task Demand	17
5	Empirical Model	18
5.1	From Correlation to Causation	18
5.2	Baseline Causal Model: Staggered DiD	18
5.3	Instrumental Variable	19
5.3.1	Relevance	19

Firm-Level Effects of AI on Labor Markets

5.3.2	The Exclusion Restriction	19
5.3.3	IV Implementation	20
6	Empirical Analysis	20
6.1	Baseline Regression Results and Analysis	20
6.2	Staggered Difference-in-Differences	20
6.3	Instrumental Variable Estimates	21
6.4	Further Impacts: Wages and New Job Creation	21
7	Discussion and Policy Implications	21
7.1	Reallocation, Not Annihilation	21
7.2	A Policy Toolkit for the Great Reallocation	22
7.3	The Road Ahead	24
7.4	Limitations and Future Research	24
8	Conclusion	25
	References	27
	Acknowledgements	28

1 Introduction

In recent years, artificial intelligence (AI)-represented by machine learning, large language models (LLM), and automation technologies-has penetrated the economic and social sectors at an unprecedented pace, reshaping the global labor market landscape. Yet, this process is inherently double-edged: AI is not only substituting repetitive labor but also spawning new occupations such as "prompt engineers" and "AI ethicists". This duality has ignited a heated and ongoing debate between policy makers, industry leaders, and scholars. However, despite considerable discussion, our understanding of the current impact of AI on the labor market remains surprisingly unclear, often fluctuating between optimistic predictions and dystopian anxieties. Furthermore, the current market is characterized by a high degree of uncertainty and a high sense of urgency and requires immediate insight into the unprecedented influence of AI on the labor market.

Significant gaps persist in existing research regarding the microlevel mechanisms through which AI specifically affects labor demand. On one hand, much of the literature focuses on macro-level effects (e.g., industry-level employment substitution rates) or single mechanisms (e.g., "substitution effect" or "creation effect"). On the other hand, AI application is often measured using superficial methods like keyword searches (e.g., "AI" or "machine learning"), which fail to capture the actual depth of AI penetration in enterprises. Coupled with unresolved endogeneity issues (e.g., high-skill enterprises may proactively adopt AI and adjust labor structures), these limitations undermine the reliability of causal inferences. Additionally, critical issues such as the generative logic of emerging occupations and the dynamic changes in skill premiums still lack empirical support from microlevel data.

Our paper constructs a novel indicators, "Employment AI Intensity (AIFE)", using large language models (LLMs). Specifically, we examine: (1) the reallocation patterns of routine/complex/creative labor; (2) the wage polarization effect driven by AI-induced skill premiums; By bridging the analytical gap between "macro trends" and "micro mechanisms," we aim to provide micro-level evidence for understanding employment transitions in the AI era, while offering scientific insights for policymakers to optimize skill training systems and guide human-AI collaborative development.

The structure of the remaining sections of the paper is as follows. Section 2 reviews the relevant literature and elaborates on the innovations of this study; Section 3 presents the theoretical framework and research hypotheses; Section 4 introduces the data sources and variable

construction; Section 5 empirically analyzes the impact of AI on the demand of enterprises for different types of occupations; and Section 6 concludes with key findings and policy recommendations.

2 Literature Review and Contribution

2.1 Technological Progress and the Labor Market

Early studies on the impact of technological progress on the labor market were primarily grounded in the Skill-Biased Technological Change (SBTC) framework (Katz and Murphy, 1992; Acemoglu and Autor, 2011). This theory posits that technological change tends to complement high-skilled labor while substituting for low-skilled labor, thereby leading to rising skill premiums and increased income inequality (Goldin and Katz, 2008). However, the SBTC theory struggles to explain the "job polarization" phenomenon observed in European and U.S. countries since the 1990s—where employment growth has occurred in high-skill and low-skill jobs, while medium-skill jobs have shrunk (Goos and Manning, 2007; Autor and Dorn, 2013).

To account for this, Routine-Biased Technological Change (RBTC) theory has gradually become the mainstream analytical framework (Autor et al., 2003; Acemoglu and Autor, 2011). RBTC shifts the unit of analysis from "skills" to "tasks," emphasizing that technology substitutes for "routine tasks" rather than labor with specific skill levels. Autor et al. (Autor et al., 2003) further categorize work tasks into five types: non-routine analytical, non-routine interactive, non-routine manual, routine cognitive, and routine manual. This classification better captures the heterogeneous effects of technology on labor demand.

2.2 The Impact of Artificial Intelligence on Labor Demand

Research on the mechanisms through which AI affects labor demand focuses on several dimensions. First, the *displacement effect*: AI reduces demand for routine labor by automating repetitive, rule-based tasks. Technologies based on machine learning and natural language processing—such as OCR, RPA, and chatbots—have already been widely applied in fields like document processing, data entry, and customer service (Frey and Osborne, 2017). Graetz and Michaels (Graetz and Michaels, 2018) found that industrial robot adoption led to significant reductions in routine manual jobs in manufacturing. In the Chinese context, Yan et al. (Yan et al., 2020) and Wang and Dong (Wang and Dong, 2020) provided similar evidence.

Second, the *productivity effect*: AI enhances enterprise productivity and reduces production costs, potentially

boosting labor demand (especially for non-routine cognitive jobs like R&D, management, and design) through scale expansion (Brynjolfsson and McAfee, 2014). Bessen (Bessen, 2019) noted that while automation substitutes for some labor, it also indirectly creates new jobs by improving total factor productivity (TFP).

Third, the *new task creation effect*: Acemoglu and Restrepo (Acemoglu and Restrepo, 2018) proposed an “automation–new task balance” model, emphasizing that technology not only substitutes old tasks but also creates new ones. Generative AI (e.g., GPT, diffusion models) is giving rise to entirely new occupational fields, such as prompt engineers, AI ethicists, and data strategists (Eloundou et al., 2023). Felten et al. (Felten et al., 2023) found that high-exposure occupations also contain substantial augmentation (rather than substitution) opportunities.

Some scholars have also highlighted the *skill restructuring and skill premium change effect*: AI application alters the structure of skill demand, driving “skill rebalancing.” On one hand, demand for hard skills (e.g., programming, data analysis, AI operations) rises; on the other hand, the value of soft skills (e.g., communication, creativity, ethical judgment) becomes increasingly prominent (Deming, 2017). The emergence of generative AI may even weaken traditional “experience premiums,” enabling low-experience workers to boost productivity through AI tools (Noy and Zhang, 2023).

2.3 Generative AI and Large Language Models

Generative AI (e.g., large language models), due to breakthroughs in natural language understanding, content generation, and complex reasoning, differs significantly from traditional AI in its impact on the labor market. Zeng et al. (Zeng et al., 2025) argue that generative AI may improve income distribution across groups—and even narrow income gaps—through a “capital–skill complementarity” effect.

Zhang and Dan (Zhang and Dan, 2025), using Chinese recruitment data, found that hiring demand for high-exposure occupations (e.g., accounting, editing, programmers) declined, while demand for low-exposure occupations (e.g., catering services, nursing) rose. Eloundou et al. (Eloundou et al., 2023) estimated that approximately 80% of U.S. workers have at least 10% of their tasks affected by large language models (LLMs), with higher exposure among high-education, high-wage occupations.

In terms of methodology, early studies relied on occupational codes (Cortés et al., 2016) or surveys (?) for classification, facing issues like strong subjectivity and neglect of intra-occupational task differences. In recent years, advancements in natural language processing (NLP)—particularly the application of LLMs—have

enabled fine-grained task classification based on recruitment texts. Chen et al. (?) used the Chinese-BERT-wwm model to classify recruitment information with a 93% accuracy rate, significantly improving the precision and scientific rigor of classification.

2.4 Heterogeneous Impacts

Regarding heterogeneous impacts, differences across enterprises, industries, and regions are key. Non-state-owned enterprises and high-tech firms are more likely to adopt AI, leading to more significant labor structure adjustments (?). State-owned enterprises, due to their greater social responsibility, exhibit weaker displacement effects of AI on employment (Yin, 2023).

Manufacturing, finance, and IT services are high-frequency AI application sectors, experiencing the most drastic changes in labor demand structure (?). In services, face-to-face interaction-intensive jobs (e.g., nursing, education) are less affected (Frey and Osborne, 2017). Among regions, developed areas—boasting better technological infrastructure and high-skilled labor agglomeration—benefit more from AI, whereas less-developed regions face risks of deepening technological divides (Han et al., 2023). The remote collaboration features of generative AI may partially mitigate regional development imbalances (Zeng et al., 2025).

While existing literature provides a solid theoretical foundation and rich empirical evidence for understanding AI’s impact on labor demand, it still has limitations: (1) Research on generative AI remains in its infancy, with systematic insights into its differentiated effects on different skill, task, and experience groups yet to be established; (2) There is a lack of fine-grained data grounded in China’s local context, with most studies relying on macro-level data or indirect matching; (3) Enterprise-level heterogeneity analysis is insufficient, particularly comparative studies across ownership types, industries, and regions; (4) Policy research has largely focused on macro-level recommendations, lacking precise intervention designs rooted in micro-level mechanisms.

2.5 Our Findings and Contributions

To address these gaps, this study presents a rigorous examination of AI’s real-world impact on the labor market. We recognize the significance of traditional economic models. Complementing our empirical analysis, we incorporate a conventional macroeconomic model to frame aggregate dynamics and inform identification. We then depart from conventional proxies and instead use a massive dataset comprising over 100 million online job postings collected from major Chinese recruitment websites between 2016 and 2020. To precisely gauge AI’s influence, we

introduce novel metrics such as AI occupation exposure AI intensity for employment (AIFE) that systematically quantify the susceptibility of specific occupations to transformation driven by AI adoption. Employing robust econometric techniques, including instrumental variable (IV) and difference-in-differences (DID) methodologies, we rigorously estimate the causal effects of AI implementation on firm-level hiring practices and wage structures.

Our findings make clear that artificial intelligence is reshaping work. AI readily takes over routine manual and cognitive tasks, yet it also amplifies the importance of roles that require analysis, creativity, and judgment. Perhaps most striking, it is spawning entirely new kinds of jobs, both in AI-focused industries and across the broader economy. The pattern is already visible. Positions such as administrative clerks and assembly-line operators are in decline, while demand is surging for data analysts, research and development specialists, and AI practitioners. As this shift unfolds, workers who bring skills that complement AI are beginning to command a premium in the labor market, signaling a broader revaluation of expertise. In this case, alarmists claiming about mass technological unemployment miss the mark. AI is less a job-destroyer than a job-reconfigurer—transforming occupations, creating new pathways for work, and opening space for richer forms of collaboration between humans and machines.

The structure of the remaining sections of the paper is as follows. Section 2 reviews the relevant literature and elaborates on the innovations of this study; Section 3 presents the theoretical framework and research hypotheses; Section 4 introduces the data sources and variable construction; Section 5 is the empirical Model; Section 6 empirically analyzes the impact of AI on the demand of enterprises for different types of occupations; and Section 7 concludes with key findings and policy recommendations.




Figure 1. The complexity of AI Technology

3 Theoretical Framework

3.1 A Functional Definition of AI

A credible empirical analysis of Artificial Intelligence (AI) demands a definition that is both analytically precise and conceptually robust. Common definitions, often centered on vague notions of machines that “learn” or “think,” are insufficient for economic inquiry, as they lack a clear, falsifiable basis. Conversely, a definition that is merely a static list of existing algorithms, such as support vector machines or convolutional neural networks, is immediately obsolete upon the next technological breakthrough. The challenge, therefore, is to define AI not by its technical components, but by its economic function.

The history of automation has been a relentless process of substituting capital for human labor in specific tasks. The technologies of the Industrial Revolution, from the power loom to the steam engine, primarily automated tasks requiring manual dexterity and physical power. In the 20th century, early computing and robotics extended this frontier to routine cognitive tasks governed by explicit, rule-based instructions, such as calculation or record-keeping (Acemoglu and Autor, 2011). For two centuries, however, a vast range of tasks involving complex pattern recognition, implicit knowledge, and adaptation remained stubbornly resistant to automation, and are thus extremely reliant on human labour.

The recent history of AI marks a fundamental break from this pattern. The rapid succession of milestones, repurposed here not as a definition but as evidence of an accelerating frontier, illustrates this shift. The success of AlexNet in 2012 demonstrated that machines could automate complex visual pattern recognition. The development of Generative Adversarial Networks (GANs) in 2014 and Transformer architectures in 2017 showed that capital could be applied to sophisticated tasks of creation and natural language understanding. The current proliferation of Large Language Models (LLMs) represents a further, dramatic expansion of this frontier into domains of communication, summarization, and reasoning. Consequently, the 2020s are the age of large language models. The GPT suite, Grok, Gemini, and Deepseek are merely several of the large variety of LLMs that exist in the status quo. The unprecedented rise of artificial intelligence (AI) in recent years has led scientists to anticipate further advances as technology matures.

Through historical timeline, evidence reveals the core economic function of modern AI: it is the technology that enables capital to perform tasks whose execution relies on statistical inference and pattern recognition, rather than pre-programmed rules. Hence, in this paper, Machine Learning (ML) and Artificial Intelligence (AI) refer not to specific computer systems and technologies. Instead, we

conceptualize AI as *the ongoing expansion of automation into a domain of tasks previously thought to be the exclusive purview of human cognition*. In essence, we frame AI not as a single invention, but as a process: a moving frontier of what is technologically possible for capital to perform.

3.2 Theoretical Analysis

As the basis of our story, we predict that, when holding all else constant, greater AI penetration simultaneously (i) reduces firms' demand for routine, easily-automated tasks and (ii) raises the productivity of cognitive, innovation-intensive activities. This dual prediction is not trivial: aggregate labor inputs and innovation rates both tend to rise over time, so we must isolate how AI “subtracts” from routine work even as it “adds” to R&D effort. To do so, we embed a single AI-penetration parameter

$$\theta \in [0, 1]$$

into a quality-ladder framework.

3.2.1 SUBSTITUTION CHANNEL

We define the *effective* routine-labor input as

$$L_R^{\text{eff}}(t) = (1 - \theta) L_R(t), \quad (1)$$

where $L_R(t)$ is the aggregate routine labor input and θ is the fraction of routine tasks automated by AI. Differentiating with respect to θ yields

$$\frac{\partial L_R^{\text{eff}}}{\partial \theta} = -L_R(t) < 0, \quad (2)$$

so a marginal increase in AI penetration directly “subtracts” from routine labor. This formalizes our *substitution effect*: as θ rises, firms require fewer workers in standardized, repetitive roles.

3.2.2 COMPLEMENTARITY CHANNEL

In each of a continuum of product lines $j \in [0, 1]$, R&D labor $L_{RD,j}(t)$ generates quality-improving innovations according to a Poisson process with arrival rate

$$z_j(t) = \chi(\theta) L_{RD,j}(t), \quad (3)$$

where $\chi(\theta)$ is strictly increasing ($\chi'(\theta) > 0$). Thus, AI acts as a complement by raising the productivity of R&D labor.

Aggregating across all lines, the *creative-destruction hazard*—the rate at which incumbents are displaced—is

$$\delta(t) = \int_0^1 z_j(t) dj = \chi(\theta) L_{RD}(t), \quad (4)$$

with $L_{RD}(t) = \int_0^1 L_{RD,j}(t) dj$.

Each successful innovation raises the quality of line j by a factor $(1 + \mu)$, so the aggregate quality index

$$A(t) = \int_0^1 A_j(t) dj$$

evolves as

$$\dot{A}(t) = \mu \int_0^1 z_j(t) A_j(t) dj = \mu \chi(\theta) L_{RD}(t) A(t). \quad (5)$$

Hence the balanced-growth rate of quality is

$$g(t) = \frac{\dot{A}(t)}{A(t)} = \mu \chi(\theta) L_{RD}(t). \quad (6)$$

Because $\chi'(\theta) > 0$, raising θ increases both the creative-destruction rate δ in (4) and the long-run growth rate g in (6). This captures our *complementarity effect*: AI not only automates routine tasks, but also “creates” demand for innovation-driving human effort.

Intuition. Equation (1) formalizes how routine tasks, which are easily codified and automated, are directly substituted by AI, reducing the demand for low-skilled labor. In contrast, equations (3)–(6) show that AI increases the marginal productivity of high-skill R&D, leading to both a higher rate of creative destruction (since incumbents are supplanted) and a faster accumulation of aggregate quality. Together, these two channels display our assumption that AI is both a destructive force for routine work and a constructive force for innovation-intensive activities.

3.2.3 JOB POLARIZATION

As AI technology becomes more deeply integrated into various industries, it does not simply restructure existing jobs. While AI penetration affects the demand for current labor, it also spurs the creation of entirely new industries and occupations, such as AI trainers, data labeling engineers, equipment maintenance technicians, etc. The incorporation of AI Capital as a newly-emerged factor of production naturally and necessarily leads to the manifestation of unique job categories linked to AI management, maintenance, and development, as a further complement to AI-driven automated systems. Moreover, due to the technological nature of AI, these occupations often require extensive specialized knowledge and a high degree of expertise, including technical skills, human judgment, and experience with supervision and management of complex systems. As a result, while demand in the labor market will increase as a whole, the newly-arisen demand will likely be concentrated in highly-specialized, creative labor rather than routine labor.

Simultaneously, the further polarization of various occupational groups leads to a rapidly widening gap in

their salary and living conditions. While the high demand for technicians and specialized, “creative” labor drives up equilibrium wages, routine labor faces a high threat of mass displacement and low market demand, translating into continuously lower wages. In other words, society becomes highly stratified, with the “high skill premium” being awarded to the emerging elite workforce, while routine labor faces wage stagnation or declination due to overabundance of supply. For workers who are unable to reskill, they are incapable of entering the high-barrier technological labor market with significantly higher wages. This dual effect manifests in a growing wage gap between occupations, exacerbating existing inequalities.

We define NC_{it} as the proportion of total tasks in the firm i at time t that are *newly created* (i.e. did not exist at the baseline t_0). We further define E_{it} as the AI exposure index for firm i at time t . Specifically,

$$NC_{it} = \beta_0 + \beta_1 \cdot E_{it} + \beta_2 X_{it} + \mu_i + \delta_t \quad (7)$$

In which X_{it} is the general skill level of the labor force, and μ_i and δ_t represent firm and time-fixed effects. $\beta_1 > 0$, indicating that higher AI exposure results in a greater share of newly created tasks.

We further define the average wage level w_{it} , with w_{it}^{high} and w_{it}^{low} representing creative and routine tasks respectively.

$$\begin{aligned} w_{it}^{\text{low}} &= \alpha_0^{\text{low}} + \alpha_1^{\text{low}} \cdot E_{it} + \alpha_2^{\text{low}} X_{it} + \mu_i + \delta_t \\ w_{it}^{\text{high}} &= \alpha_0^{\text{high}} + \alpha_1^{\text{high}} \cdot E_{it} + \alpha_2^{\text{high}} X_{it} + \mu_i + \delta_t \end{aligned}$$

The fact that $\alpha_1^{\text{low}} < 0$ and $\alpha_1^{\text{high}} > 0$ indicates that while AI penetration has a positive influence on wages of creative labor, the opposite is true for routine labor. This completes the empirical proof for our core assumptions.

3.3 A Task-Based Model of AI and Labor

While our ultimate goal is an empirical quantification of AI’s impact, a purely empirical approach risks measurement without theory. To guide our investigation and provide a disciplined framework for interpreting our results, we first develop a dynamic task-based model of the firm. This section lays the theoretical groundwork for the subsequent empirical analysis, recognizing that any model is a simplification, but one that is indispensable for clarifying the mechanisms at play. We build upon the canonical framework of Acemoglu and Restrepo (2018), but innovate by introducing firm-level heterogeneity in AI capacity and endogenizing the creation of new tasks at the firm level. This structure allows us to derive a rich set of testable predictions

about firm-level hiring and wage structures, directly linking our theory to our empirical strategy.

3.3.1 THE ECONOMIC ENVIRONMENT

We model the production side of the economy by focusing on a representative firm, indexed by f . The firm’s decisions are the micro-level engine of the phenomena we wish to study.

Production Technology and Tasks. The firm produces a final good, Y_f , by performing a continuum of tasks indexed by $i \in [N_t - 1, N_t]$. This task space is normalized to have a measure of one, but its frontier, N_t , can expand over time through innovation. Following the standard literature, we assume these tasks are combined with a constant elasticity of substitution, σ :

$$Y_f = \left(\int_{N_t-1}^{N_t} y_f(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}, \quad (8)$$

where $y_f(i)$ is the quantity of output from task i . We assume $\sigma > 1$, implying that tasks are substitutes.

Factors of Production and Task Allocation. We depart from the single-labor-type model and assume the firm can hire two distinct types of labor, distinguished by the nature of the tasks they perform.

- **Routine Labor (L_R):** Performs standardized, codifiable tasks.
- **Creative Labor (L_C):** Performs complex, non-routine tasks, which include problem-solving, management, and, crucially, research and development (R&D).¹

In addition to labor, the firm can employ **AI-Capital (K)**.

The allocation of these factors to tasks is governed by two thresholds. A threshold J divides tasks that are inherently routine from those that are inherently complex. The key decision variable for the firm is the **automation threshold**, $I_f \in [N_t - 1, J]$.

- Tasks $i \in [N_t - 1, I_f]$ are **automated** and can be produced only by AI-capital.
- Tasks $i \in (I_f, J]$ are **routine** and can be produced only by Routine labor.

¹We combine the performance of existing complex tasks and the creation of new tasks into a single labor category. This abstracts from potential skill heterogeneity within the high-skill workforce, a simplification that would grant us a clearer focus on the firm’s allocation between routine and creative functions as a whole.

- Tasks $i \in (J, N_t]$ are **complex** and can be produced only by Creative labor.

For analytical tractability, we treat J as an exogenous parameter reflecting the fundamental technological possibilities of an era, while I_f is a choice variable for the firm.

It is important to note that this task-allocation structure is a deliberate simplification from canonical models (e.g., Acemoglu and Restrepo (2018)), where capital and labor often compete to perform the same task based on relative costs. Our approach, which assigns factors to distinct task domains, is chosen for tractability and to cleanly isolate the impact of firm-specific AI capacity, θ_f , on the two key margins of interest: the automation margin (I_f) and the new task creation margin ($\dot{N}_t$). This focuses the analysis on how a firm's core capability drives its technological choices, rather than on substitution effects from short-run factor price fluctuations.

Firm-Specific AI Capacity. The central innovation of our model is the introduction of a firm-specific parameter, $\theta_f \in [0, \bar{\theta}]$, representing the firm's **AI capacity**. This parameter reflects the firm's accumulated knowledge, infrastructure, and organizational capital for deploying AI. It is not technology itself, but the firm's capability to leverage it. This capacity has a dual effect on productivity.²

The output of a task i is linear in the assigned factor, with productivity $\gamma(\cdot)$:

$$y_f(i) = \begin{cases} \gamma_K(\theta_f)k_f(i) & \text{if } i \leq I_f \\ \gamma_R l_{R,f}(i) & \text{if } I_f < i \leq J \\ \gamma_C l_{C,f}(i) & \text{if } J < i \leq N_t \end{cases} \quad (9)$$

where $k_f(i)$, $l_{R,f}(i)$, and $l_{C,f}(i)$ are the factor quantities. To isolate the role of AI, we assume γ_R and γ_C are constant. The productivity of AI-capital, however, depends on the firm's capacity: we assume $\gamma'_K(\theta_f) > 0$. A firm with higher AI capacity is more efficient at using AI-capital to perform automated tasks.

The Dynamics of Task Creation. The task space is not static. Firms can create new, more complex tasks through R&D, which is performed by Creative labor. Let $L_{C,f}^{R\&D}$ be the amount of Creative labor allocated to R&D. The frontier of tasks N_t evolves according to:

$$\dot{N}_t = \psi(\theta_f) L_{C,f}^{R\&D}. \quad (10)$$

²For tractability, we treat θ_f as an exogenous parameter. Endogenizing its accumulation, for instance through learning-by-doing or as a function of R&D investment, is a compelling aspect for future research but beyond the scope of this paper's objective to derive static predictions.

Here, $\psi(\theta_f)$ is the productivity of R&D, and we make the crucial assumption that $\psi'(\theta_f) > 0$. This formalizes the **complementarity channel**: a firm with greater AI capacity is more effective at innovation. AI tools for data analysis, simulation, and discovery augment the abilities of its creative workforce. New tasks created at the frontier N_t are, by their nature, complex and initially fall into the domain of Creative labor, generating what the literature calls a **reinstatement effect**.

3.3.2 THE FIRM'S OPTIMIZATION PROBLEM

The firm chooses its automation threshold I_f , its allocation of factors to production, and its R&D investment to maximize the present discounted value of its profit stream. It takes factor prices as given: w_R for Routine labor, w_C for Creative labor, and r for AI-capital.

Static Cost Minimization and the Automation Decision.

At any point in time t , for a given state N_t , the firm first solves a static problem of minimizing the cost of producing a given amount of output Y_f . This requires choosing the optimal intensity of each task, $y_f(i)$. The price of one unit of output from task i , denoted $p_f(i)$, is equal to its marginal cost of production:

$$p_f(i) = \begin{cases} r/\gamma_K(\theta_f) & \text{if } i \leq I_f \\ w_R/\gamma_R & \text{if } I_f < i \leq J \\ w_C/\gamma_C & \text{if } J < i \leq N_t \end{cases} \quad (11)$$

The firm's choice of the automation threshold I_f is governed by a trade-off. To automate a task i , the firm must incur a one-time, sunk cost, which we assume is a smooth, decreasing function of the task index, $C(i)$, with $C'(i) < 0$. This captures the idea that simpler, lower-indexed tasks are cheaper to automate. The firm will automate all tasks up to the point I_f where the marginal benefit of automating one more task equals the marginal cost. The benefit is the perpetual stream of cost savings from using cheaper AI-capital instead of Routine labor.

Proposition 3.1 (The Optimal Automation Threshold). *The profit-maximizing automation threshold I_f is determined by the condition where the marginal cost of automation equals the present value of the flow of cost savings:*

$$\begin{aligned} C(I_f) &= \int_0^\infty e^{-\rho t} \left(\frac{w_R}{\gamma_R} - \frac{r}{\gamma_K(\theta_f)} \right) dt \\ &= \frac{1}{\rho} \left(\frac{w_R}{\gamma_R} - \frac{r}{\gamma_K(\theta_f)} \right) \end{aligned} \quad (12)$$

where ρ is the firm's discount rate. This equation implicitly defines I_f as a function of factor prices and the firm's AI capacity, $I_f(\theta_f, w_R, r)$.

From this condition, we can determine how the firm's automation strategy responds to its AI capacity. An increase in θ_f raises $\gamma_K(\theta_f)$, which lowers the cost of using AI-capital and thus increases the cost-saving from automation on the right-hand side of (12). To restore equality, the firm must move to a higher I_f where the marginal cost of automation, $C(I_f)$, is lower (since $C'(i) < 0$). This gives us our first key result.

Corollary 3.2 (The Displacement Channel). *The optimal automation threshold I_f is strictly increasing in the firm's AI capacity θ_f :*

$$\frac{\partial I_f}{\partial \theta_f} > 0. \quad (13)$$

This is the mathematical formalization of the displacement effect. A higher AI capacity directly incentivizes the firm to expand the range of tasks performed by machines, displacing Routine labor.

The Dynamic R&D Decision. The firm's dynamic problem is to choose the amount of Creative labor to allocate to R&D, $L_{C,f}^{R\&D}$, to control the evolution of its task frontier N_t . Let $V_f(N_t)$ be the firm's value function. The associated Hamilton-Jacobi-Bellman (HJB) equation is:

$$\rho V_f(N_t) = \max_{L_{C,f}^{R\&D}} \left\{ \Pi_f(N_t, L_{C,f}^{R\&D}) + \dot{N}_t \frac{\partial V_f(N_t)}{\partial N_t} \right\}, \quad (14)$$

where Π_f is the instantaneous profit flow, given by total revenue minus total factor costs (including the wages of R&D workers). Substituting $\dot{N}_t$ from (10), the first-order condition with respect to $L_{C,f}^{R\&D}$ is:

$$\frac{\partial \Pi_f}{\partial L_{C,f}^{R\&D}} + \psi(\theta_f) \frac{\partial V_f(N_t)}{\partial N_t} = 0. \quad (15)$$

The marginal cost of hiring one more R&D worker is their wage, w_C . The marginal benefit is the value created by accelerating the expansion of the task frontier. Thus, the optimality condition simplifies to a beautifully intuitive expression:

$$w_C = \psi(\theta_f) \frac{\partial V_f(N_t)}{\partial N_t}. \quad (16)$$

The term $\frac{\partial V_f(N_t)}{\partial N_t}$ is the shadow value of a new task. Equation (16) states that the firm hires Creative labor for R&D up to the point where their wage equals their marginal value product in innovation. This value product is magnified by the firm's AI capacity, θ_f .

3.3.3 FROM MODEL TO MEASUREMENT: DERIVING TESTABLE HYPOTHESES

We now have all the necessary components to derive the model's predictions for firm-level labor demand. These predictions will form the basis of our empirical work, where

our constructed measure of AI intensity will serve as a proxy for the theoretical parameter θ_f .

Let's first write down the firm's demand for each type of production labor. Integrating the standard CES demand for each task's output over the relevant range gives us the demand for each factor:

$$L_{R,f} = \int_{I_f(\theta_f)}^J l_{R,f}(i) di = (J - I_f(\theta_f)) \cdot \left(\frac{P\gamma_R}{w_R} \right)^\sigma Y_f \quad (17)$$

$$L_{C,f}^{Prod} = \int_J^{N_t} l_{C,f}(i) di = (N_t - J) \cdot \left(\frac{P\gamma_C}{w_C} \right)^\sigma Y_f \quad (18)$$

Total demand for Creative labor is $L_{C,f} = L_{C,f}^{Prod} + L_{C,f}^{R\&D}$.

Hypothesis 1 (Substitution): *Higher firm-level AI capacity (θ_f) leads to a decrease in the firm's demand for Routine labor (L_R).*

Derivation: We can directly differentiate the expression for $L_{R,f}$ with respect to θ_f . The primary channel of impact is through the automation threshold $I_f(\theta_f)$.

$$\frac{\partial L_{R,f}}{\partial \theta_f} \propto -\frac{\partial I_f}{\partial \theta_f}. \quad (19)$$

From Corollary 13, we proved that $\frac{\partial I_f}{\partial \theta_f} > 0$. Therefore, it follows unambiguously that $\frac{\partial L_{R,f}}{\partial \theta_f} < 0$. The model predicts that firms with greater AI capacity will structurally reduce their demand for workers in routine tasks. This is not an assumption, but a result of the firm's optimization over the automation margin.

Hypothesis 2 (Complementarity): *Higher firm-level AI capacity (θ_f) leads to an increase in the firm's demand for Creative labor (L_C).*

Derivation: The effect on Creative labor is driven by the R&D channel. From the optimality condition (16), for a given shadow value of new tasks, an increase in θ_f raises the productivity of R&D, $\psi(\theta_f)$. To maintain the equality, the firm must increase its hiring of R&D labor, $L_{C,f}^{R\&D}$, which in turn accelerates the creation of new tasks, $\dot{N}_t$. These new tasks are performed by Creative labor, subsequently increasing $L_{C,f}^{Prod}$ over time. Both effects are positive.

$$\frac{\partial L_{C,f}}{\partial \theta_f} = \underbrace{\frac{\partial L_{C,f}^{R\&D}}{\partial \theta_f}}_{>0 \text{ (Direct R\&D Effect)}} + \underbrace{\frac{\partial L_{C,f}^{Prod}}{\partial N_t} \frac{\partial N_t}{\partial \theta_f}}_{>0 \text{ (Reinstatement Effect)}} > 0. \quad (20)$$

The model thus generates a powerful and robust complementarity between a firm's AI capacity and its demand for high-skilled, creative workers.

Hypothesis 3 (Wage Polarization): *The skill premium, reflected in the relative demand for Creative versus Routine labor, increases with a firm’s AI capacity (θ_f).*

Derivation: While our partial equilibrium model takes wages as given, we can analyze the firm’s relative labor demand, which is the force that would drive wage changes in a general equilibrium setting. We examine the derivative of the ratio of demand for Creative to Routine labor:

$$\frac{\partial}{\partial \theta_f} \left(\frac{L_{C,f}}{L_{R,f}} \right) = \frac{\frac{\partial L_{C,f}}{\partial \theta_f} L_{R,f} - L_{C,f} \frac{\partial L_{R,f}}{\partial \theta_f}}{(L_{R,f})^2}. \quad (21)$$

The numerator is unambiguously positive, as we have shown that $\frac{\partial L_{C,f}}{\partial \theta_f} > 0$ and $\frac{\partial L_{R,f}}{\partial \theta_f} < 0$. A higher AI capacity systematically skews a firm’s labor demand structure towards Creative labor. This provides a clear theoretical basis for expecting to find a higher wage premium for AI-complementary skills within firms that are more intensive AI adopters. While our partial equilibrium model takes wages as given, this predicted shift in the firm’s relative labor demand is precisely the micro-level force that, when aggregated, would be expected to drive wage polarization in a general equilibrium setting.

3.4 Numerical Simulation

To translate the theoretical mechanisms of our model into quantitatively meaningful predictions, we conduct a numerical simulation. While the analytical results in the previous section established the qualitative directions of AI’s impact, a calibrated simulation allows us to explore the magnitude of these effects and visualize the full general equilibrium reallocation of economic activity. The objective is to trace how the economy’s balanced growth path—including the allocation of labor, the rate of innovation, and the pace of creative destruction—evolves as AI penetration, θ , increases from a baseline of zero to a state of high integration.

3.4.1 CALIBRATION

To ground our simulation in plausible economic reality, we calibrate the model’s parameters. Our strategy is to use standard values from the macroeconomic and endogenous growth literature for established parameters, while choosing reasonable functional forms and values for the novel parameters governing AI’s impact. The complete set of parameter values used for our baseline simulation is presented in Table 1.

The total labor supply, L_{total} , is normalized to one. The step size of each quality-improving innovation, μ , is set to 0.2, a value consistent with models of endogenous technological

change. The core of our simulation lies in the specification of the R&D productivity function, $\chi(\theta)$. We adopt a simple and transparent linear form: $\chi(\theta) = \chi_0(1 + \gamma\theta)$. The baseline R&D productivity, χ_0 , is set to 0.1. The most critical parameter is γ , which governs the elasticity of R&D productivity with respect to AI penetration. In our baseline specification, we set $\gamma = 2.0$. We explore the sensitivity of our results to this crucial parameter in a subsequent analysis.

Table 1. Baseline Parameter Calibration.

PARAMETER	VALUE	DESCRIPTION
L_{TOTAL}	1.0	TOTAL LABOR SUPPLY
μ	0.2	INNOVATION STEP SIZE
χ_0	0.1	BASILINE R&D PRODUCTIVITY
γ	2.0	AI COMPLEMENTARITY ELASTICITY

Notes: This table lists the parameter values used in the baseline numerical simulation. The functional form for R&D productivity is $\chi(\theta) = \chi_0(1 + \gamma\theta)$.

3.4.2 SIMULATION RESULTS

Using the calibration in Table 1, we solve for the model’s balanced growth path equilibrium for each value of θ in the range $[0, 0.9]$. We present the results sequentially to build the economic narrative step by step.

First, we examine the most direct consequence of rising AI penetration: the reallocation of the workforce. Figure 2 plots the equilibrium shares of labor allocated to routine and R&D tasks. At $\theta = 0$, with no AI, the economy dedicates the vast majority of its workforce (98.0%) to routine tasks, with only a small fraction (2.0%) engaged in R&D. As AI penetration increases, the relative productivity of R&D rises, pulling workers out of the routine sector. This reallocation is initially slow but accelerates as θ becomes larger. By the time AI has automated 90% of routine tasks ($\theta = 0.9$), the labor allocation has dramatically shifted: the share of the workforce in R&D has surged to 35.9%, while the share in routine tasks has fallen to 64.1%. This demonstrates that AI is a potent driver of structural transformation in the labor market.

This reallocation is driven by two powerful, competing forces, which we visualize in Figure 3. The solid red line (left axis) shows the *effective* routine labor input, $L_R^{\text{eff}} = (1 - \theta)L_R$. This quantity declines sharply with AI penetration, driven by both the direct substitution effect (the $1 - \theta$ term) and the indirect reallocation effect (the decline in L_R). At $\theta = 0.9$, the effective routine labor input has fallen to just 0.064, less than 7% of its initial level. This is the powerful displacement force of AI. Simultaneously, the dashed blue line (right axis) shows the R&D productivity, $\chi(\theta)$. This represents the powerful complementarity force,




Figure 2. Labor Reallocation. This figure plots the equilibrium allocation of labor between routine tasks (L_R) and R&D (L_{RD}) as a function of AI penetration (θ).

making the remaining innovators in the economy vastly more effective.




Figure 3. Substitution vs. Complementarity. The solid red line (left axis) shows the effective routine labor input. The dashed blue line (right axis) shows the productivity of R&D labor. The two opposing forces of AI are shown as a function of AI penetration (θ).

The central question for macroeconomic performance is which of these two forces dominates. Figure 4 shows the net effect on the economy's long-run growth rate, g . The result is unambiguous: the complementarity effect ultimately overwhelms the substitution effect in driving aggregate performance. As AI penetration increases, the growth rate accelerates, initially slowly and then exponentially. At $\theta = 0$, the economy is nearly stagnant, with a growth rate of just 0.04%. However, at $\theta = 0.9$, fueled by both a larger

R&D workforce and higher R&D productivity, the long-run growth rate surges to 2.01% per year. This result provides a clear, quantitative illustration of our model's prediction that AI, despite its displacement effects, can be a powerful engine of long-run prosperity.




Figure 4. Long-Run Economic Growth. This figure plots the economy's balanced-growth-path annual growth rate as a function of AI penetration (θ).

Finally, we consider the impact on economic dynamism, measured by the rate of creative destruction, δ . As shown in Figure 5, this rate, which measures the pace at which new innovations displace old technologies, follows a path similar to the growth rate. It rises from a near-zero base to a rate of 0.10 at $\theta = 0.9$. This implies that in a high-AI economy, the expected lifespan of a given technology or firm is significantly shorter. While this heightened churn is a necessary component of faster growth, it also suggests a more volatile economic environment for firms and workers.

3.4.3 SENSITIVITY ANALYSIS

The magnitude of the complementarity effect is governed by the parameter γ . To ensure our optimistic conclusion about long-run growth is not merely an artifact of our baseline calibration, we conduct a sensitivity analysis. We re-run the simulation under two alternative scenarios: a "Low Complementarity" case where $\gamma = 1.0$, and a "High Complementarity" case where $\gamma = 3.0$.

The results are qualitatively robust but quantitatively informative. In the Low Complementarity scenario ($\gamma = 1.0$), the long-run growth rate at $\theta = 0.9$ is a more modest 1.05%. In the High Complementarity scenario ($\gamma = 3.0$), the growth rate reaches a remarkable 3.15%. This analysis confirms that while the qualitative result of AI-driven growth acceleration is a robust feature of the model, the ultimate magnitude of this effect is highly dependent on the degree




Figure 5. Creative Destruction. This figure plots the rate of creative destruction (δ) as a function of AI penetration (θ).

to which AI technologies augment human ingenuity. This highlights the importance of policies and firm strategies that foster human-AI collaboration in innovation as a key determinant of long-run economic performance.

4 Sample Data Selection and Research Design

Our empirical investigation is built upon a vast, granular dataset of online job postings, which we process using a state-of-the-art NLP pipeline. This section details the construction of our analytical corpus. We first introduce the source of our data, the Web Data Commons, and outline its unique advantages over traditional labor market surveys. We then describe, in detail, the complete process of extraction, normalization, and cleaning, which utilizes Large Language Models (LLMs) to create a robust, firm-level panel dataset from noisy, web-scale structured data. Our work transforms unstructured text into a novel metric, the AI Intensity for Employment (AIFE), required for the causal analysis that follows.

4.1 The Job Postings Universe

Traditional analyses of labor market dynamics have long been constrained by the frequency and granularity of official government statistics. Surveys like the Bureau of Labor Statistics' Occupational Employment Statistics (OES) are typically released with a significant lag and present data at an aggregate, macroeconomic level. This, in turn, obscures the high-frequency firm-level adjustments that are central to understanding the impact of rapid technological change such as AI.

To overcome these limitations, we turn to the internet for a much more specific approach: online job postings. Our data




Figure 6. Our Data Treatment Process. This process ensures we obtain an unique and novel result, departing from conventional proxies.

is sourced from the Web Data Commons (WDC) project, a major academic initiative that extracts structured data from the Common Crawl. The Common Crawl refers to a collaborative and authoritative project launched and developed by major search engines, including Google, Bing, Yahoo, etc. that aims to maintain a standardized vocabulary for online data. It is also one of the largest and most comprehensive public web crawls available (2016). Specifically, we utilize the WDC's Schema.org extraction, focusing on data annotated with the `JobPosting` type. The use of the `schema.org` standard, while not universally or perfectly adopted, provides large-scale structured data extracted from company websites worldwide, offering firm-level statistics and observations in numerous states and industries globally.

The raw data set comprises all publicly available WDC `JobPosting` extractions spanning from January 2015 to December 2024, encompassing several major tech booms of the 21st century. The initial dataset consists of approximately 412 million data quads, which is a vast but extremely noisy and heterogeneous collection of job advertisements. The immediate challenge then lies in transforming this raw material into a consistent, research-grade panel data set.

After selection, we extract relevant information from the dataset, including:

1. **Firm identifiers and location:** Unique firm IDs, country codes, and other information that allow firms to be linked across years and enables aggregating results

by region.

2. **Employment-related information:** online job postings, job descriptions, special requirements and skills listed in job postings, usually contained in the url.




Figure 7. JobPosting property coverage (Circa 2022 Release).

By combining these firm-level characteristics with employment and wage information, we can critically examine how the presence and penetration of AI correlate with labor market alterations.




Figure 8. (a) Annual volume of unique job postings, 2015-2024. Note that the 2024 statistics is currently available.

After extraction, the primary advantage of this corpus is obvious:

1. **Scale and Scope:** It offers a near-census of online job postings, dwarfing proprietary datasets from single platforms and providing a more holistic view of labor demand.
2. **Timeliness:** The data is available at long range, from 2015 to 2024, meaning that we are equipped with a tool that tracks labor market changes through a significant time period.

However, the data's web-scraped nature presents significant challenges, including rampant duplication, lack of sufficient and reliable information, and missing or malformed fields. The remainder of this section details the rigorous pipeline we developed to address these issues.

Data Processing and Normalization Pipeline. To convert the raw job postings data into an analytical corpus, we implemented a multistage pipeline designed for hectogigabyte-scale data.

1. *Extraction and Encoding Unification:* The first step involves parsing the raw data quads and unifying disparate field names (e.g., 'name' vs. 'title') into a consistent schema for key attributes: `job_title` and `job_description`. Garbled text, identification sequences, and non-natural language encodings are removed for clarity purposes. Postings with unreadable paths or queries are discarded.

2. *LLM-Powered Field Normalization:* A central challenge is the extremely low amount of information included in a url string. An example url may appear as "https://xxx.xxx.com/jobs/1/", this gives little information and thus have to be discarded in this second roll. Similarly, urls only containing vague phrases like "jobs-on-post" must be eliminated in this process. Rule-based methods (e.g., regular expressions) are too brittle to handle this variety.

To solve this, we employ a Large Language Model for structured data extraction. The model was trained to perform two tasks: (i) extract the canonical job name from the `job_title` field, stripping away meaningless suffixes and extraneous information. And (ii) generate a description concerning aspects of the job.

3. *Firm Name Canonicalization:* While the LLM provides a clean baseline, further steps are needed to create unique firm identifiers for our panel analysis. We apply a second layer of canonicalization to the LLM-cleaned firm names. This involves: (a) converting all names to a uniform case and encoding; (b) applying fuzzy string matching algorithms to join highly similar names (e.g., "finances" and "financing").

The output of this pipeline is a clean, deduplicated corpus unique job postings from over 13 years. Table 2 provides summary statistics of the final analytical dataset.

4.2 Measuring the Independent Variable: A Firm-Level AI Intensity Index

To measure our key theoretical construct of firm-level AI capacity (θ_f), we introduce our novel metrics: AI intensity for employment (AIFE). We depart from simple keyword semantic search that traditional studies usually employ. Instead, we use a Gemma-family LLM to perform a

Table 2. Summary Statistics of the Analytical Job Posting Corpus (2015-2024).

STATISTIC	VALUE
TIME PERIOD	2015 - 2024
TOTAL RAW POSTINGS (QUADS)	~412 MILLION
POSTINGS AFTER PARSING	~203 MILLION
POSTINGS AFTER DE DUPLICATION	~158 MILLION
UNIQUE FIRMS IDENTIFIED	~1.2 MILLION
UNIQUE JOB TITLES (RAW)	~4.1 MILLION
MEDIAN POSTINGS PER FIRM-YEAR	18




Figure 9. (b) Annual website volume 2015-2024.

context-aware assessment of each job posting's relationship to AI.

Definition 4.1 (AI Intensity for Employment, AIFE). The AIFE is our primary continuous measure. For firm f in year t , it is the average AI score across all jobs $j \in J_{ft}$ posted by that firm. This captures the overall AI-centricity of a firm's hiring strategy.

$$\text{AIFE}_{ft} = \frac{1}{|J_{ft}|} \sum_{j \in J_{ft}} \text{Score}_j / 10. \quad (22)$$

We prompted the LLM to rate each job description, specifically by assigning a score from 0 to 10 reflecting its depth of AI integration. The rubric distinguishes between:

- **Core AI Roles (Score 8-10):** Involve the research and development of new AI models or systems (e.g., ML Researcher).
- **Applied AI Roles (Score 4-7):** Involve the application of existing AI tools to solve business problems (e.g., Data Scientist using AutoML).
- **AI-Aware Roles (Score 1-3):** Use AI-powered software but do not require direct AI expertise (e.g., Marketing Analyst using an AI-powered suite).

LLM Regression Prompt

You are an expert data analyst specialized in evaluating the impact of artificial intelligence on the labor market. Your task is to assess the "AI Penetration" of a given job occupation. Definition of AI Penetration: The degree to which the core responsibilities of this job currently involve or require the use of Artificial Intelligence (AI) and Machine Learning (ML) technologies. Instructions: Analyze the job title provided. Classification: Determine if the job is fundamentally AI-related. If it is not, output exactly 0. Scoring: If the job is AI-related, assign it a single integer score from 1 to 10 based on the following criteria: 9=8-10: Core AI Development Research. Jobs focused on creating new AI algorithms, models, and systems (e.g., "AI Researcher," "Machine Learning Engineer"). 4-7: AI Application Integration. Jobs that require significant expertise to apply, manage, or maintain existing AI systems to solve complex problems or enhance workflows (e.g., "Data Scientist," "Business Intelligence Developer"). 1-3: AI Tool Assistance. Jobs that primarily use AI-powered tools and software as a component of their broader, non-AI-specific tasks (e.g., "Digital Marketer using analytics tools," "Writer using a grammar checker with AI"). Output Format: Your response must be only a single number: 0 or a score between 1 and 10. Do not include any other text, explanations, or labels. The job to analyze: sample

The classification of jobs is a highly intuitive process, yet we proposed self-improvement methods to reduce the hallucinative bias. The classification LLM is forced to give specific reasons and reflect upon its decision. It is also encouraged to reduce potential bias by exploring its past explanations. This rigorous process yields our primary dependent variable: the share of a firm's new hires in a given year that fall into the 'Routine' category, and similarly for 'Creative' and 'Complex' roles. This allows us to directly test how a firm's AIFE score affects the composition of its workforce.

4.3 Measuring Dependent Variables: The Task Content of Labor Demand

We then categorize jobs into three types based on the canonical task-based literature: Creative, Complex, and Routine. This directly finds statistical evidence for our core hypothesis.

The same Gemma-family LLM was employed here for this large-scale classification task. For each job description, the model was prompted to perform a zero-shot classification, assigning the job to one of the three categories based on detailed definitions provided in the prompt. The prompt structure is illustrated below.

The identity of a job analyst was first assigned to the LLM. Routine occupations were defined by their "dependence on predictable, manual, or cognitive tasks that follow a well-defined set of rules". Complex jobs are defined as "multifaceted roles requiring advanced cognitive skills, adaptive problem-solving in dynamic environments, and the management of numerous interdependent tasks and stakeholders." Creative jobs are defined by its "focuses on generating original ideas, solving problems innovatively, and producing progressive content". The complete prompt is shown as follows:

LLM Classification Prompt

Role: You are an expert literary analyst, skilled in deconstructing themes, characters, and narrative structures. You will now apply your analytical skills to the modern "text" of a job description.

Task: Analyze the provided job title: "job". Classify its primary nature based on the following definitions: Creative: Generating original ideas, solving problems innovatively, and producing novel content. Complex: Involves multifaceted roles requiring advanced cognitive skills, adaptive problem-solving in dynamic environments, and the management of numerous interdependent tasks and stakeholders. Routine: Characterized by a dependence on predictable, manual, or cognitive tasks that follow a well-defined set of rules.

Output Instructions: 1. First, state your classification as either "Creative," "Complex," or "Routine." 2. Then, provide a detailed explanation in the form of a brief analytical essay. Justify your classification by dissecting the job title against the definitions provided. Explore the nuances and potential ambiguities. 3. Finally, include a brief reflection on the societal or economic context of this

role. Why might this classification matter?

Style: Write with the precision and insight of a critic analyzing a seminal novel. Avoid any markdown or code. Your response should be purely prose.




Figure 10. Evolution of the Task Composition of Hiring, 2015-2024.

4.4 Statistical Findings

Table 2 presents our annual firm-level employment shares for three occupational groups over 2015 to 2024. By construction, the category shares sum to 100% each year. The composition at the endpoints is:

- 2015: Creative 16.8%, Complex 48.7%, Routine 34.5%.
- 2024: Creative 31.9%, Complex 50.8%, Routine 17.3%.

Two features are immediately apparent. First, there is a large decline in Routine work, falling by 27.0 percentage points (pp). Second, the lost Routine share is primarily reallocated to Creative (+18.5 pp) and, to a lesser extent, Complex (+8.5 pp).

The adjustment is not linear. From 2015 to 2019, Routine drops from 34.5% to 23.8% while Complex rises to 55.1% and Creative advances modestly to 21.1%. In 2020, due to disturbance on logistics and employment conditions mainly attributable to the COVID-19 Pandemic, the composition tilts sharply toward Complex (51.2%) with a temporary trough in Creative (24.3%) and a slight Routine uptick (24.5%). From 2021 onward the reallocation accelerates away from Routine with Creative growing quickly (23.4% → 31.9% by 2024). Complex retrenches in 2022–2023

(54.9% → 52.1%) and partially rebounds in 2024 (57.2%). Thus, the post-2020 period is characterized by a pronounced Routine → Creative substitution and a milder Routine → Complex shift.

A compositional perspective sharpens these patterns. The log ratio

$$\ln \left(\frac{\text{Creative}}{\text{Routine}} \right)$$

rises from -0.72 in 2015 to 1.55 in 2024, implying Creative work moves from roughly $0.49\times$ Routine to about $4.7\times$ Routine by the end of the sample. Similarly,

$$\ln \left(\frac{\text{Complex}}{\text{Routine}} \right)$$

increases from 0.34 to 2.03 —from $1.4\times$ to $7.6\times$ Routine. These transformations remove the constant-sum constraint and indicate sustained substitution away from Routine rather than a mere artifact of normalization.

Rank dynamics tells the same story. Complex remains the largest category throughout (rank 1). Routine holds rank 2 through 2020 but falls to rank 3 in 2021 and stays there as Creative overtakes it and continues to widen the gap through 2024.

Table 3. Labor Composition Shares by Year (2015–2024)

Year	Creative	Complex	Routine
2015	0.1682	0.4866	0.3452
2016	0.1792	0.5000	0.3208
2017	0.1713	0.5281	0.3002
2018	0.1876	0.4931	0.3193
2019	0.2111	0.5510	0.2379
2020	0.2432	0.5122	0.2446
2021	0.2335	0.6128	0.1537
2022	0.2771	0.5492	0.1737
2023	0.2845	0.5380	0.1775
2024	0.3185	0.5081	0.1734

4.5 The Correlation Between AI and Task Demand

Moving beyond simply describing the data, one can obtain both intuitive and implicit correlation between metrics. In addition, a credible causal claim first requires an underlying correlation. There are two entrenched correlations observed in our data: the sharp acceleration of AI adoption in corporate hiring and the simultaneous, dramatic reconfiguration of the task content of labor demand.

The Post-2017 Acceleration of AI Adoption. While the intellectual roots of artificial intelligence are deep, the technology’s transition from a niche academic pursuit to




Figure 11. (b) Labor Force Composition.

a transformative general-purpose technology is a recent phenomenon. We identify 2017 as a critical inflection point. The publication of the seminal paper “Attention Is All You Need” in that year introduced the Transformer architecture, a novel neural network design that proved exceptionally effective at handling sequential data (2017). Due to the fact that the architecture replaced sequential bottlenecks with massively parallel training and made long-range dependencies easy to model, this breakthrough revolutionized the field of natural language processing. More importantly, it also laid the architectural foundation for the current generation of Large Language Models (LLMs) and generative AI. Its impact was to dramatically lower the industry-level cost and increase the feasibility of applying AI to a vast new range of cognitive tasks previously unable to be achieved by AI algorithms.

Our data reflects this technological shock. As shown in the left panel of Figure 12, the average AI Intensity for Employment (AIFE) score across all job postings exhibits a sharp and sustained increase, particularly after 2017. Simultaneously, the right panel shows a dramatic reconfiguration of the labor market: the share of routine jobs collapses, while the share of creative jobs nearly doubles. These aggregate patterns are the macroeconomic reflection of the firm-level dynamics we seek to understand.

The Reconfiguration of Labor Tasks. Contemporaneous with this surge in AI hiring, our data reveals a profound shift in the composition of the workforce. Figure 12 and Table 3 document the evolution of hiring shares for Creative, Complex, and Routine occupations. The most dramatic trend is the precipitous decline in the share of Routine work, which falls from over a third of all job postings in 2015 to just 17.3% by 2024. This collapse is followed by a significant rise in the share of Creative work, which nearly doubles over the same period.




Figure 12. Aggregate Trends in AI Intensity and Task Composition (2015-2024). The left panel plots the average AIFE score across all job postings. The right panel plots the aggregate share of Routine and Creative jobs.

These aggregate trends are suggestive, but they do not, on their own, link the decline of routine work to the rise of AI. The two trends could be driven by separate, unrelated factors. The crucial next step is to move from aggregate time-series to the firm level and examine whether the firms driving the AI boom are the same firms driving this labor market reconfiguration.

5 Empirical Model

5.1 From Correlation to Causation

In order to ascertain the causal relationship between AI incorporation and industry labor demand, it is critical to systematically eliminate potential endogeneity issues that could interfere with causal inference. This paper utilizes the following methods to remove endogeneity.

5.2 Baseline Causal Model: Staggered DiD

We build a panel regression model as our causal model. We first build a simplified two-period DiD that provides a convenient baseline representation of the average treatment effect for intuition. Specifically, we define:

$$Y_{it} = \beta_0 + \beta_1 \text{Post}_t + \beta_2 \text{Treat}_i + \beta_3 (\text{Treat}_i \times \text{Post}_t) + \gamma' \mathbf{X}_{it} + \varepsilon_{it}, \quad (23)$$

where $\text{Treat}_i \in \{0, 1\}$ indicates treated firms and Post_t is the post-treatment period.

We further develop our validation by constructing a staggered DiD regression model. Because AI adoption is unevenly distributed across firms and the treatment effect may vary inconsistently with time, we group firms into cohorts based on their adoption date. For each cohort, we utilize a relative-time event study to model how the treatment evolves relative to the adoption date. We thereby eradicate the disruption to the final result brought by varying lengths of AI adoption for different firms at any specific instant in time. We then aggregate cohort-specific effects into overall treatment effects using modern staggered DiD estimators, which addresses biases that arise when adoption timing is heterogeneous. Consequently, this model resolves the issue of varying AI adoption times for different industries.

$$Y_{it} = \alpha_i + \eta_t + \gamma' \mathbf{X}_{it} + \varepsilon_{it} + \sum_{g \in \mathcal{G}} \sum_{s \neq -1} \beta_{g,s} \mathbf{1}\{G_i = g\} \cdot \mathbf{1}\{t - t_0 = s\}, \quad (24)$$

where:

- G_i denotes the cohort when firm i first adopts AI (or $G_i = \infty$ if never treated),
- g denotes the specific cohort for individual firms
- s is relative time from adoption (we omit the relative period $s = -1$ as the reference category).
- t_0 refers to the adoption time ($t_0 + s$ refers to the standard time)
- $\beta_{g,s}$ refers to the average treatment effect coefficient for cohort g at relative time s

For each cohort g and standard time t we define the group-time average treatment effect as:

$$\text{ATT}(g, t) = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid G_i = g, t \geq g]. \quad (25)$$

An aggregation of the combined effect at relative time s can then be formed by weighted addition of cohort-specific effects:

$$\hat{\beta}_s = \sum_{g: t_0 + s \in \mathcal{T}} w_{g,s} \widehat{\text{ATT}}(g, t_0 + s), \quad (26)$$

where $w_{g,s}$ are weights proportional to cohort size, and $\mathcal{T}$ is the set of observed calendar periods.

We present two complementary representations of the staggered DiD model. Equation (23) is a cohort-by-relative-time event-study specification that estimates cohort-specific dynamic effects $\beta_{g,s}$ for each adoption cohort g at relative time s . Equation (24) defines the group-time ATT for cohort g at standard time t , and Equation (26) shows how these group-time ATTs are aggregated into combined effect $\hat{\beta}_s$.

In order to implement this model, we need to validate several critical assumptions. The adoption of staggered treatment is satisfied, as once firms systematically receive treatment, they will likely continue and further expand AI incorporation throughout the observation period, rather than returning to the previous state. AI often bring systematic, structural changes that are difficult to revoke.

The Parallel Trends assumption is also satisfied. The parallel trends assumption requires that, in the absence of treatment, the outcome for the treatment and control groups would have evolved in the same way, i.e. the difference between the two groups would have remained constant. Formally, for $\forall t > 0$,

$$E[Y_{it}(0) | D_i = 1] - E[Y_{it}(0) | D_i = 0] = \Delta_0 \quad (27)$$

In which $Y_{it}(0)$ refers to the outcome without treatment, D_i refers to treatment on either group, and Δ_0 refers to the time-independent difference in outcome. The assumption is satisfied, as AI adoption can be viewed as exogenous to short-term labor market trends. While the **incentive** for AI incorporation primarily results from an industry-level or state-level macroeconomic push, the **capacity** for AI incorporation (e.g., the necessary capital and technology required for specific firms) is primarily caused by the exposure to a firm-level shock. Therefore, it could be reasonably assumed that within an industry, despite varying degrees of actual implementation, the difference in general labor demand for firms would have been negligible were the treatment not present due to homogeneous incentives.

The significance of such a model is that it not only demonstrates the magnitude of the treatment effect, but also resolves the inaccuracies brought by standard, two-period DiD when treatments (AI adoption) occur at varying times across industries with different characteristics (and thus different responses to an AI introduction). The traditional DiD framework functions effectively only when analyzing scenarios containing only two periods and a consistent treatment effect in between. In this model, adopters at numerous different time periods can be accurately compared using non-treated units as controls, making staggered DiD

the more realistic model for analyzing AI's real-world implications to the economy.

5.3 Instrumental Variable

To complement our DiD estimates and further strengthen our causal claims, we employ an Instrumental Variable (IV) strategy. The primary endogeneity concern is that a firm's decision to adopt AI may be correlated with unobserved, time-varying factors that also affect its hiring patterns (e.g., a new management team, unobserved demand shocks). To address this, we construct a "shift-share" or "Bartik-style" instrument for a firm's AI intensity that is standard in the modern labor economics literature (2018).

The instrument is designed to predict a firm's exposure to AI based on two components: (1) its historical, pre-treatment occupational hiring patterns (the "share"), and (2) the subsequent national-level growth in AI within each occupation (the "shift"). Formally, we construct the predicted AI exposure for firm i in year t as:

$$\text{Predicted_AIFE}_{it} = \sum_k (\text{Share}_{k,i,t_0} \times \text{National_AI_Growth}_{k,t}), \quad (28)$$

where Share_{k,i,t_0} is the share of occupation k in firm i 's total hiring during a pre-treatment base period (t_0 , e.g., 2012-2015), and $\text{National_AI_Growth}_{k,t}$ is the growth in AI-related job postings for occupation k at the national level between t_0 and t , excluding firm i 's own postings to avoid mechanical correlation.

5.3.1 RELEVANCE

The relevance of this instrument relies on the plausible assumption that firms with a historical concentration in occupations that later experienced a national-level surge in AI adoption will, on average, see a greater increase in their own AI intensity. For example, a firm that heavily employed "data analysts" and "statisticians" in 2014 is more likely to be an intensive AI adopter in 2020 than a firm that primarily employed "manual assemblers," because the former occupations were at the epicenter of the national AI boom. We will formally test this in our first-stage regression, where we expect our `Predicted_AIFE` instrument to be a strong and statistically significant predictor of the firm's actual AIFE score. A first-stage F-statistic well above the conventional threshold of 10 will provide evidence against a weak instrument problem.

5.3.2 THE EXCLUSION RESTRICTION

The validity of our IV strategy hinges on the exclusion restriction: the instrument, `Predicted_AIFE`, must affect a firm's hiring composition *only through* its effect on the firm's actual AI adoption. This requires that

the national-level AI trends within occupations are not correlated with firm-specific shocks to labor demand, conditional on our controls.

We argue this is plausible. The national trend in AI adoption for, say, accountants, is driven by the aggregate actions of thousands of firms, technological breakthroughs from major labs, and evolving industry standards. It is unlikely that these broad, national forces are directly correlated with the idiosyncratic, short-run demand shocks experienced by a single accounting firm in a specific city, especially after we control for firm fixed effects (which absorb all time-invariant firm characteristics) and industry-by-year fixed effects (which absorb common shocks to a firm's specific industry). The instrument's identifying power comes from the fact that the firm's pre-period occupational shares are historical and pre-determined, making them plausibly exogenous to the shocks of the AI-boom era. We provide empirical support for this assumption with a placebo test in the results section.

5.3.3 IV IMPLEMENTATION

We implement the IV strategy using a two-stage least squares (2SLS) framework.

First Stage: We regress the firm's actual AI intensity on our shift-share instrument and the full set of controls. This stage isolates the portion of a firm's AI adoption that is plausibly exogenous, as it is driven by the interaction of its historical structure and national trends.

$$\text{AIFE}_{it} = \delta_0 + \delta_1 \text{Predicted_AIFE}_{it} + \zeta' \mathbf{X}_{it} + \mu_i + \eta_t + \nu_{it}, \quad (29)$$

Second Stage: We then regress the outcome variable (e.g., the share of routine jobs) on the predicted values of AI intensity, $\widehat{\text{AIFE}}_{it}$, obtained from the first stage.

$$Y_{it}^k = \beta_0 + \beta_1^{\text{IV}} \widehat{\text{AIFE}}_{it} + \gamma' \mathbf{X}_{it} + \mu_i + \eta_t + \varepsilon_{it}^k \quad (30)$$

where Y_{it}^k is the share of labor of type k for firm i in year t , ε_{it}^k is our AI intensity measure, $\mathbf{X}_{it}$ are time-varying controls, μ_i are firm fixed effects, and η_t are year fixed effects.

6 Empirical Analysis

6.1 Baseline Regression Results and Analysis

We begin our empirical results section by establishing the baseline conditional correlations between AI intensity and key labor market outcomes. This analysis proceeds in three stages. First, we examine the aggregate time-series trends in our data. Then, we use panel regression models to test whether these correlations hold after controlling for

a set of observable firm characteristics and unobserved heterogeneity. While these models are not designed to yield causal estimates, they serve as a crucial first step in validating our hypotheses and motivating the causal analysis that follows.

We now turn to a more rigorous test using firm-level panel data. Our baseline specification is a fixed effects model of the following form:

$$Y_{ft}^k = \beta_0 + \beta_1 \cdot \text{AIFE}_{ft} + \gamma' \mathbf{X}_{ft} + \mu_f + \eta_t + \varepsilon_{ft}, \quad (31)$$

where Y_{ft}^k is the share of labor of type k for firm f in year t , AIFE_{ft} is our AI intensity measure, $\mathbf{X}_{ft}$ are time-varying controls, μ_f are firm fixed effects, and η_t are year fixed effects.

Evidence on Wages: The AI Skill Premium. We explore the relationship between AI intensity and wages, a key component of our third hypothesis on wage polarization. Table 4 presents results from regressions where the dependent variable is the average wage offered for jobs in each task category. Columns 1-3 show the raw correlations, while Columns 4-6 include year and job posting fixed effects.

The results are striking. Across all specifications, AI intensity is positively and significantly associated with wages for all job types. However, the magnitude of this "AI premium" is highly heterogeneous. In the full specification (Columns 4-6), the coefficient on AI is largest for Routine jobs, followed by Creative and then Complex jobs. This suggests that as firms adopt AI, they not only shift their hiring away from routine roles but also increase the skill requirements—and thus the wages—for the routine jobs that remain. The positive and significant coefficient for Creative and Complex jobs is consistent with a complementarity effect, where AI adoption raises the productivity and pay of high-skilled workers.

Together, these baseline results establish a powerful and robust set of conditional correlations that are consistent with our theoretical framework. They provide a strong empirical foundation for the causal models that follow, which are designed to address the remaining issue of potential endogeneity.

6.2 Staggered Difference-in-Differences

We now present our primary causal estimates from the staggered DiD model specified in Equation (24). This model identifies the effect of AI adoption by comparing changes in hiring patterns within firms after they adopt AI to the contemporaneous changes in firms that have not yet adopted.

Table 5 presents the full regression results, averaging the post-treatment effects. Column 1 shows the effect on the

Firm-Level Effects of AI on Labor Markets

Table 4. Baseline Regression Results: Association between AI Intensity and Wages by Job Category

	(1)	(2)	(3)	(4)	(5)	(6)
Variables	Routine	Creative	Complex	Routine	Creative	Complex
AI Intensity (AIFE)	-0.251*** (0.032)	0.189*** (0.021)	0.057* (0.033)	-0.105*** (0.023)	0.082*** (0.011)	0.021* (0.036)
Control Variables	YES	YES	YES	YES	YES	YES
Year FE	NO	NO	NO	YES	YES	YES
Industry FE	NO	NO	NO	YES	YES	YES
State FE	NO	NO	NO	YES	YES	YES
Year-State FE	NO	NO	NO	YES	YES	YES
Year-Industry FE	NO	NO	NO	YES	YES	YES
Constant	-0.0376 (1.3e-5)	0.0228 (7.9e-6)	0.030 (2.8e-5)	-0.164 (6.9e-6)	0.207 (4.8e-5)	0.233 (2.0e-5)
Observations	2.43e4	4.08e4	1.89e4	2.43e4	4.08e4	1.89e4
Adj. R-squared	0.062	0.139	0.091	0.164	0.146	0.033

Notes: The table reports coefficients from OLS regressions where the dependent variable is the average wage. The unit of observation is a job posting. T-statistics are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

share of routine jobs, confirming the visual evidence from the event study. The average post-treatment effect is a statistically significant reduction of [e.g., 5.2] percentage points. Column 2 repeats the analysis with the share of creative jobs as the outcome, finding a positive and significant increase of [e.g., 3.1] percentage points. Column 3 shows a smaller but still positive effect on complex jobs. Columns 4-6 demonstrate that these results are robust to the inclusion of a rich set of time-varying firm-level controls, such as firm size (total postings) and industry-by-year fixed effects.

6.3 Instrumental Variable Estimates

To strengthen our causal claim, we employ an IV strategy to identify our claim. Table 6 presents the 2SLS results. The first three columns shows the first-stage regression, confirming that our instrument is to some extent relevant to AIFE, but not significantly relevant to the job count. Routine-task workers were significantly less AI-penetrated, supporting our hypothesis of AI's substitution effect. Creative-task workers were also significantly more AI-penetrated, supporting our hypothesis of AI's complementary aspect.

The latter three columns present results after controlling for year, industry, and state variables. The estimated effect of AI adoption on the routine share is -0.127, remarkably consistent with our DiD estimate and our hypothesis of the substitution effect. The effect on the creative share is also positive and significant. This consistency across two distinct identification strategies, which rely on different assumptions, provides strong evidence that our main findings are not an

artifact of a single methodological choice.

6.4 Further Impacts: Wages and New Job Creation

Finally, we explore the impact of AI adoption on wages and the creation of new types of jobs. The first two columns of table 7 present estimates where the outcome is the average log wage offered in a firm's job postings. We find that AI adoption is associated with a significant increase in the average offered wage (0.071), consistent with a skill-biased shift in labor. This also confirms our third hypothesis about wage polarization. The latter two columns shows that AI-adopting industries are significantly more likely to post for novel occupations, providing direct evidence for the job creation channel.

7 Discussion and Policy Implications

7.1 Reallocation, Not Annihilation

As we previously mentioned, the debate over the future development of artificial intelligence is long cast in a false dichotomy, with opinions ranging from a jobless future to a utopian paradise. This has created a policy vacuum whereby no action is taken because of the alleged uncertainty. The evidence in this research points towards a reality that is more complex and immensely urgent. A nineteenth-century way of thought is not going to define the economic landscape of the twenty-first century; rather, the twenty-first century will demand a workforce with the broadest spectrum of skills. Our results—reflecting an almost annihilatory downward shift in demand for routine work from over a third of all jobs to an almost inexorable decline in less than a decade, with

Firm-Level Effects of AI on Labor Markets

Table 5. The Impact of AI Adoption on the Task Content of Hiring: Staggered DiD Estimates.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	ROUTINE	CREATIVE	COMPLEX	ROUTINE	CREATIVE	COMPLEX
AI-ADOPTION (β)	-0.0528*** (0.004)	0.0735*** (0.008)	0.1198*** (0.025)	-0.049*** (0.008)	0.029*** (0.006)	0.020* (0.011)
CONTROL VARIABLES	YES	YES	YES	YES	YES	YES
YEAR FE	YES	YES	YES	YES	YES	YES
INDUSTRY FE	YES	YES	YES	YES	YES	YES
STATE FE	YES	YES	YES	YES	YES	YES
OBSERVATIONS	2.43E4	4.08E4	1.89E4	2.43E4	4.08E4	1.89E4
R-SQUARED	0.016	0.029	0.036	0.173	0.224	0.149

Notes: The table reports average post-treatment coefficients from the staggered DiD estimation (Equation (24)). The dependent variable is the share of the indicated job type in a firm's hiring. 'Post-Adoption' is an indicator variable equal to one for all periods after a firm adopts AI. Firm controls include [e.g., log total postings, firm age]. Standard errors, clustered at the firm level, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6. Endogeneity: Instrumental Variable

Variables	(1)	(2)	(3)	(4)	(5)	(6)
	Routine	Creative	Complex	Routine	Creative	Complex
AI Intensity (AIFE)	-0.237*** (0.041)	0.077*** (0.052)	0.028* (0.033)	-0.127*** (0.033)	0.057*** (0.004)	-0.031* (0.046)
Control Variables	YES	YES	YES	YES	YES	YES
Year FE	NO	NO	NO	YES	YES	YES
Industry FE	NO	NO	NO	YES	YES	YES
State FE	NO	NO	NO	YES	YES	YES
Constant	-0.0218 (1.2e-5)	0.0513 (8.1e-6)	0.042 (3.2e-5)	-0.128 (6.9e-6)	0.516 (4.7e-5)	0.211 (2.2e-5)
Observations	2.43e4	4.08e4	1.89e4	2.43e4	4.08e4	1.89e4
Adj. R-squared	0.048	0.227	0.021	0.331	0.395	0.147

Notes: The table reports coefficients from OLS regressions where the dependent variable is the average wage. The unit of observation is a job posting. T-statistics are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

that disorder complemented by a rise in demand for jobs that demand creative or complex skills—are not a forecast. They are in a direct lateral transformation well underway.

To speak of "robots taking jobs" here is to fundamentally misread the situation. A more accurate metaphor is a powerful economic tide that is rapidly eroding the bedrock of routine-based employment while simultaneously creating new high ground for those with analytical and creative skills. The policy challenge, therefore, is not to build sea walls in a futile attempt to hold back the tide, but to equip the global workforce with the means to navigate these turbulent waters and reach higher ground. A failure to do so will not result in mass joblessness, but in something perhaps dangerous to a democratic society: a permanent and unbridgeable chasm between the AI-augmented and the AI-displaced.

7.2 A Policy Toolkit for the Great Reallocation

Confronting this "Great Reallocation" requires moving beyond macro-level anxiety and toward a portfolio of micro-level, targeted policies. The evidence demands a new policy toolkit designed not for a jobless future, but for a future of constant job transition. The old twentieth-century systems, built for a stable world of lifelong careers, are no longer fit for purpose. A new framework must be built on three pillars: a radical overhaul of skills development, a modernized social safety net that supports transitions, and a strategic focus on augmenting, rather than simply automating, human labor.

First, the world must build infrastructures for continuous reskilling. The dramatic shift away from routine cognitive tasks documented in our findings implies that the current

Firm-Level Effects of AI on Labor Markets

Table 7. Further Impacts: Wages and New Job Creation

VARIABLES	(1)	(2)	(3)	(4)
	LG(AVG.) CREATIVE WAGE	LG(AVG.) COMPLEX WAGE	% OF CREATIVE NEW JOB	% OF COMPLEX NEW JOB
AI	0.071*** (0.006)	0.024*** (0.005)	0.073*** (0.003)	0.031*** (0.001)
CONTROL VARIABLES	YES	YES	YES	YES
YEAR FE	YES	YES	YES	YES
INDUSTRY FE	YES	YES	YES	YES
STATE FE	YES	YES	YES	YES
OBSERVATIONS	2.43E4	4.08E4	2.43E4	2.43E4
R-SQUARED	0.101	0.122	0.201	0.152

Notes: The dependent variable is the share of the indicated job type in a firm's hiring. 'Post-Adoption' is an indicator variable equal to one for all periods after a firm adopts AI. Firm controls include [e.g., log total postings, firm age]. Standard errors, clustered at the firm level, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.




Figure 13. Comparison between Individuals with Different Skills. Our results indicate that individuals with analytical and innovative capabilities are more equipped to adapt to the AI-dominated economy than individuals limited to routine occupations.

model of education, where skills are acquired in the first two decades of life, is obsolete. The traditional four-year university degree, while still valuable, must be complemented by a more agile and accessible learning ecosystem. This requires a global effort to elevate and integrate community colleges, vocational programs, and certified "coding boot camps" into a unified credentialing system. Student aid programs should be reformed to support shorter, modular, and industry-recognized credentials, allowing a factory worker in Ohio to retrain as a robotics maintenance technician with the same level of federal support as a student enrolling in a liberal arts college. Furthermore, public-private partnerships should be incentivized through the tax code to create "lifelong learning accounts," co-funded by employers and the government, that empower workers to proactively acquire new skills

throughout their careers.

Second, the social safety net must be redesigned to act as a springboard, not just a cushion. The current unemployment insurance system is a blunt instrument designed for temporary layoffs during recessions. It is entirely inadequate for workers whose skills have been permanently obsolete by technology. A twenty-first century system should embrace the concept of "wage insurance," which would partially compensate workers who take a new, lower-paying job after being displaced, thereby easing their transition without creating a disincentive to reenter the workforce. More ambitiously, global governments should pilot "Transition Assistance Accounts," a flexible, government-seeded fund that a displaced worker could use not only for income support but also for relocation expenses to move to a region with better job prospects, for childcare during retraining, or as a subsidy to a small business to cover the cost of on-the-job training. Such a system would allow displaced workers to act as active agents in their own economic transition.

Finally, innovation policy must steer the trajectory of AI development. A reasonable national government must not emphasize in its ability to deploy capital, but also validate the ingenuity of its human capital. Yet, the current incentive structure often prioritizes automation that simply replaces labor over augmentation that enhances it. This is a strategic error. Governments should restructure R&D tax credits to explicitly favor the development of "human-complementary" AI—systems designed to work alongside humans to increase their productivity and creative capacity. By creating a thumb on the scale for augmentation, we can guide technological development in a direction that leverages the world's greatest asset. A nation that successfully pairs a creative workforce with augmenting AI

will out-innovate and outperform a nation that merely seeks to replace its workers with automated systems. Managing this transition effectively is not just a matter of domestic social policy; it is a geopolitical imperative.




Figure 14. Overview of Government Policy Implementation.

7.3 The Road Ahead

The political obstacles to such a vision are formidable. They cut across entrenched interests, ideological divides, and jurisdictional boundaries. Yet, the evidence of a profound structural shift is now undeniable. To ignore it is to risk a future of deepening inequality, social fragmentation, and a decline in national competitiveness. The challenge is not technological; it is political.

What is required is a new national compact between government, industry, and the worker. Industry must move beyond a narrow focus on shareholder value to embrace its role in investing in the long-term skills of its workforce. Government must move beyond its twentieth-century bureaucratic silos to create a flexible, modern system that empowers workers to adapt. And the people must embrace a culture of lifelong learning as a cornerstone of economic security.

Thinking of the advance of Artificial Intelligence as a wave that washed away resistance is one of those outdated thoughts that we used to have. It may be true that the numbers speak louder than words that this re-shaping of the labor market has picked up pace. The great challenge that this generation of policymakers faces is to prepare institutions that withstand shocks and a workforce that can be quickly adapted so that individual countries and the world may prosper in the landscape of change. Therefore, the great challenge of this generation is, not to halt AI, but to build the institutional and educational scaffolding that allows human talent to flourish alongside it.

7.4 Limitations and Future Research

While our study provides one of the first large-scale, firm-level analyses of AI's impact on the labor market



Figure 15. A conceptual framework outlining the core themes for analyzing AI's impact on labor demand

using a novel measurement strategy, its conclusions must be understood within the context of its methodological boundaries.

First, our analysis is based on the universe of online job postings, which captures the *inflow* of labor demand—a firm's hiring decisions. It does not, however, directly observe labor outflows (such as layoffs or retirements) or, perhaps more importantly, the internal reallocation of tasks among existing employees. It is entirely possible that the effects we document are even more pronounced within firms, as incumbent workers are retrained and their roles are redefined without a new job ever being posted. In fact, this effect is one of the cornerstones of our core hypothesis. A complete picture of AI-driven workforce transformation would require merging our hiring data with administrative records on internal human resources data. Such a combined dataset would allow researchers to disentangle the extensive margin (hiring and firing) from the intensive margin (the changing nature of existing jobs), providing a truly holistic view of firm-level adjustment.

Second, our AI intensity and task content measures, while rather methodologically advanced, are, in the end, proxies. The AIFE score captures a revealed preference for AI-related skills by firms in their hiring procedures, which represents a criterion by which one may assess and grant technological adoption but does not directly measure either investment in AI capital or the expected productivity gains obtained through this investment. The next step would be to link our AIFE measure to firm-level balance sheet and production data so that a direct test of the entire mechanism implied by our theoretical model, from AI adoption to labor market re-organization and, eventually, to its final impact on firm performance, profitability, and market share, could be assessed.

Finally, our analysis is still at a partial equilibrium with a firm-level focus. Some of the changes in the firms that are adopting AI are being documented by us. Labor is undergoing massive reallocation away from routine tasks; indeed, firms doing this will place inordinate

pressure on economy-wide wage structures and returns to education, along with economic activity in terms of geography. Furthermore, our results are placed in a certain national and institutional context; hence, the speed and nature of this transition could accelerate or slow down in an economy that differs in labor market regulations, educational systems, or compositions of industries. An integrative step consists of including our firm-level causal estimations into a quantitative general equilibrium model.

These limitations notwithstanding, our study overall provides a granular, methodologically robust foundation for understanding the firm-level mechanics of AI-driven labor market change. By moving beyond aggregate proxies and developing a new toolkit for measuring technological penetration from the ground up, we offer a clear empirical benchmark and a set of stylized facts that can guide both policy and future academic research in this critical domain.

8 Conclusion

The main contribution of our paper is to develop and implement an original methodology to measure AI adoption at the firm level, together with its consequences on task content in labor demand. We generate a novel AI intensity (AIFE) measure that captures the complex and contextual nature of AI integration by using the text of millions of job adverts and analyzing them with a state-of-the-art Large Language Model. To transit from correlation to causality, we use a twin identification strategy that combines an event-study-type staggered Difference-in-Differences specification with an Instrumental Variable strategy. The robustness of the results across our empirical strategies gives great weight to our inferences.

Our findings vividly illustrate a clear and consistent narrative of reconfiguration, not annihilation. We uncover strong causal evidence indicating that the adoption of AI systematically replaces routine tasks. At the same time, we observe that AI serves as a potent complement to non-routine work, prompting firms to substantially increase their demand for creative and complex roles. This structural transformation in labor demand leads to a discernible "AI skill premium," evidenced by higher wages for jobs that align with the new technology. Importantly, we also discover direct evidence of a path to job creation, as companies that embrace AI are significantly more inclined to hire for entirely new occupational categories. We argue that these facts presented in this document require the attention of policy makers. The current concern about a "jobless future" appears to be a misunderstanding of the root issue. The immediate and critical issue at hand is not a quantity of jobs, but a substantial and rapid realignment of their nature and quality. Our results show that the world is in the early stages of a "Great Reallocation" of work, a transformation

as fundamental as that from industry to agriculture.

Just such a transformation will necessarily encounter opposition. Unless there is a collective policy initiative, gains from this AI revolution—fast tracking of growth and stimulated productivity—may accrue in a privileged fraction of the population, fueling inequality and social polarization. The ultimate direction of our AI-driven economy will not be dictated by what clever algorithms we possess, but by how smart our institutions are and how flexible people are. The challenge of our era is not to stop AI, but to build the institutional and educational scaffolding that allows human talent to flourish alongside it.

References

- Katz, L. F., Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics*, 107(1), 35–78.
- Acemoglu, D., Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics* (Vol. 4B, pp. 1043–1171). Elsevier.
- Goldin, C., Katz, L. F. (2008). *The Race between Education and Technology*. Harvard University Press.
- Goos, M., Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *American Economic Review*, 97(2), 118–122.
- Autor, D. H., Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5), 1553–1597.
- Autor, D. H., Levy, F., Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- Frey, C. B., Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
- Graetz, G., Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753–768.
- Brynjolfsson, E., McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton.
- Bessen, J. E. (2019). AI and jobs: The role of demand. *NBER Working Paper* No. 24235.
- Acemoglu, D., Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542.
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4), 1593–1640.
- Cortés, G. M., Jaimovich, N., Nekarda, C. J., Siu, H. E. (2016). Reallocation across occupations and the decline of middle-skill employment. *Finance and Economics Discussion Series 2016-0001*, Board of Governors of the Federal Reserve System.
- Acemoglu, D., Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- Dauth, W., Findeisen, S., Südekum, J., Woessner, N. (2017). German robots – The impact of industrial robots on workers. *IZA Discussion Paper* No. 10484.
- Yan, X., Zhu, K., Ma, C. (2020). Employment under robot impact: Evidence from China’s manufacturing industry. *Statistical Research* (), 37(11), 40–52. (In Chinese)
- Wang, J., Dong, B. (2020). Industrial robots and employment: Evidence from China. *China Economic Review*, 62, 101–123.
- Han, J., Li, X., Zhang, Y. (2023). Artificial intelligence, regional development, and the geography of new work in China. *Journal of Economic Geography*, 23(6), 1161–1193.
- Yin, Z. (2023). State ownership and labor adjustment under automation: Evidence from China’s SOEs. *China Economic Quarterly*, 22(3), 97–128. (In Chinese)
- Zhang, H., Dan, Q. (2025). Generative AI exposure and firm hiring: Evidence from Chinese online recruitment. *Economic Research Journal*, 60(2), 88–106. (In Chinese)
- 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. *arXiv:2303.10130*.
- Felten, E. W., Raj, M., Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. *arXiv:2304.03843*.
- Noy, S., Zhang, W. (2023). Experimental evidence on the productivity effects of generative AI. *SSRN Working Paper* 4375283.
- Zeng, J., Liu, Y., Li, X. (2025). Generative AI, capital–skill complementarity, and income distribution. *Economic Modelling*, 132, 106581.
- Krizhevsky, A., Sutskever, I., Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097–1105.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27, 2672–2680.
- Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
- Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

- Ho, J., Jain, A., Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33, 6840–6851.
- Achiam, J., Adler, S., Amini, A., et al. (2023). GPT-4 technical report. *arXiv:2303.08774*.
- Gemini Team. (2024). Gemma: Open models based on Gemini research and technology. *arXiv:2403.08295*.
- Brinkmann, A., Heist, N., Lehmborg, O., Bizer, C. (2023). The Web Data Commons Schema.org data set series. In *Proceedings of the Web Conference Companion (WWW '23)*, 999–1007.
- Peeters, R., Ibrahim, A., Heist, N., Bizer, C. (2024). The Web Data Commons Schema.org table corpora. In *Proceedings of the Web Conference Companion (WWW '24)*, 119–126.
- Lehmborg, O., Ritze, D., Meusel, R., Bizer, C. (2016). A large public corpus of web tables containing time and context metadata. In *Proceedings of the 25th International World Wide Web Conference Companion (WWW '16)*, 75–76.
- Manku, G. S., Jain, A., Sarma, A. D. (2007). Detecting near-duplicates for web crawling. In *Proceedings of the 16th International World Wide Web Conference (WWW '07)*, 141–150.
- Charikar, M. (2002). Similarity estimation techniques from rounding algorithms. *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, 380–388.
- Cui, Y., Che, W., Liu, T., Qin, B., Yang, Z. (2019). Pre-training with whole word masking for Chinese BERT. *arXiv:1906.08101*.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. *Stanford University Working Paper*.
- Arntz, M., Gregory, T., Zierahn, U. (2016). The risk of automation for jobs in OECD countries: A comparative analysis. *OECD Social, Employment and Migration Working Papers*, No. 189.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30.
- Brynjolfsson, E., Li, D., Raymond, L. (2023). Generative AI at work. *NBER Working Paper No. 31161*.
- Bloom, N., Brynjolfsson, E., Djankov, S., McGowan, D. (2024). AI and productivity: Experimental evidence from customer support. *American Economic Review Insights*, 6(2), 123–140.
- Acemoglu, D., Lelarge, C., Restrepo, P. (2022). Competing with robots: Firm-level evidence from France. *American Economic Review*, 112(4), 1108–1145.
- Felten, E. W., Raj, M., Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *SSRN Working Paper 3368605*.
- Zhang, S., Balog, K., Adafre, S. F. (2020). Web table extraction, retrieval and augmentation: A survey. *arXiv:2002.00207*.
- Lehmborg, O., Ritze, D., Meusel, R., Bizer, C. (2017). Stitching web tables for improving matching quality. *Proceedings of the VLDB Endowment*, 10(11), 1502–1513.
- Meusel, R., Bizer, C., Lehmborg, O. (2014). The WebDataCommons Microdata, RDFa and Microformat dataset series. In *Lecture Notes in Computer Science (Vol. 8796)*, 277–292. Springer.
- Felten, E. W., Raj, M., Seamans, R. (2021). How susceptible are jobs to AI? Using AI occupational exposure to analyze labor-demand impacts. *AEJ: Insights*, 3(4), 31–48.
- Goos, M., Manning, A., Salomons, A. (2009). Job polarization in Europe. *American Economic Review*, 99(2), 58–63.
- Autor, D. H. (2019). Work of the past, work of the future. *AEA Papers and Proceedings*, 109, 1–32.
- OpenAI. (2023). Model evaluations for extreme risks. *OpenAI Technical Report*.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2018). Bartik Instruments: What, When, Why, and How. *NBER Working Paper No. 24408*.

Acknowledgements

This project represents the culmination of a collective journey to explore one of the most urgent economic questions of our time: the impact of Artificial Intelligence on the labor market. What began as a spark of shared curiosity about generative AI soon became a months-long endeavor that tested our resilience, sharpened our analytical skills, and deepened our appreciation for the power of teamwork. We were motivated by a paradox we observed in public discourse: while the rise of AI was widely debated, discussions often oscillated between utopian optimism and dystopian fear, rarely grounded in detailed, firm-level evidence. The existing literature on automation, though rich, tended to rely on aggregate proxies that overlooked the nuanced adjustments happening within companies. From this gap emerged our central research ambition: to provide a large-scale, microdata-driven analysis of how AI adoption influences hiring strategies, labor demand, and the balance between substitution, complementarity, and job creation. Our work would not have been possible without the invaluable guidance of our supervising instructors, Professors Mingmao Deng and Jingping Chen. Professor Deng played a pivotal role from the very beginning, helping us refine our research question to be both ambitious and tractable. His feedback on our theoretical models ensured they were not only innovative but firmly grounded in economic principles. His advice on empirical strategies—particularly the use of Difference-in-Differences and Instrumental Variable methods—was crucial in strengthening the rigor of our analysis. Professor Chen provided meticulous revisions to the mathematical logic of our theoretical framework and elevated the clarity of our academic writing. Though neither professor was involved in data processing or manuscript drafting, their intellectual generosity and selfless dedication shaped every stage of this work. We are profoundly grateful for their mentorship, offered without compensation, and for the example they set as scholars committed to truth and rigor. Each member of our team contributed uniquely and indispensably. As team captain, Liqian Yan guided the overall direction of the project and authored the theoretical and policy-oriented sections. Liqian developed the dynamic task-based model at the heart of our paper, derived its core propositions, and translated the empirical findings into policy implications, providing a vision for how governments might navigate the future of work in the age of AI. Yintong Chen led the literature review and the design of our empirical strategy. His systematic review of existing research positioned our work within the broader scholarly conversation, while his development of the staggered Difference-in-Differences model and the robust shift-share Instrumental Variable framework gave our analysis its methodological strength. Zichun Qiu laid the foundation of the study through rigorous data work. He managed the processing of the gigabyte-scale Web Data Commons dataset, engineered pipelines for data cleansing, and applied an LLM to extract firm names and job classifications. His construction of the AIFE indicator and the classification of job tasks—routine, complex, and creative—were critical innovations that enabled our empirical analysis. Our research journey was far from smooth, and the challenges we encountered shaped both the paper and ourselves. Data acquisition and cleaning posed the first major hurdle: the Web Data Commons dataset, though vast, was noisy and unstructured. Early rule-based methods failed, forcing us to embrace a cutting-edge solution—fine-tuning an LLM for entity normalization. Later, concerns about endogeneity demanded that we abandon simpler regression models in favor of more sophisticated strategies, including a defensible shift-share IV design. Implementing a modern staggered Difference-in-Differences estimator required careful attention to parallel trends and event study validation. Each obstacle forced us to confront the limits of our methods, but overcoming them strengthened our analysis and our resolve. Computational scale was another trial. Classifying millions

Firm-Level Effects of AI on Labor Markets

of job postings threatened to overwhelm our resources. We designed a hybrid pipeline that combined rule-based parsers for straightforward cases with LLM-based classification for complex ones, executed in parallel and batched to maximize efficiency. What initially felt like a barrier became an opportunity to innovate, leaving us with not only results but also methodological tools we hope will serve future research. Yet beyond the technical hurdles, what truly defined this project was the spirit of collaboration. Each of us brought different strengths—whether in theory, data, or methods—but what sustained us was our mutual trust, late-night problem-solving sessions, and the knowledge that every challenge was being faced together. The countless hours we spent debating equations, debugging code, and interpreting results became more than academic labor; they became a shared experience of growth, resilience, and camaraderie. The completion of this paper is therefore not only an academic milestone but also a personal one. We are indebted to Professors Deng and Chen for modeling what it means to pursue knowledge with humility and persistence. We are thankful to our families for their quiet but unwavering support, for the encouragement that carried us through long nights and moments of doubt. And above all, we are grateful to one another. As teammates who have faced obstacles side by side, we have learned that true scholarship is never a solitary pursuit—it is forged in dialogue, in collaboration, and in the willingness to carry one another forward. This project has given us more than findings and results; it has given us a glimpse of what rigorous, collaborative research can achieve. The lessons we have learned—the value of precision, the courage to confront uncertainty, and the spirit of teamwork—will remain with us as we continue our academic journeys. Whatever paths we pursue, we carry forward the conviction that knowledge, when pursued with integrity and shared effort, can illuminate even the most complex questions of our time.