Article

The Substitution Effect of Artificial Intelligence on Middle-Skill Jobs: A Core Source of Labor Market Polarization

Zhichen Wan [1,*]

[1] Hainan Vocational University of Science and Technology, Haikou, China
* Correspondence: Zhichen Wan, Hainan Vocational University of Science and Technology, Haikou, China

Abstract: The rapid integration of artificial intelligence into workplace settings has intensified concerns about its asymmetric impact on occupational structures. This study investigates how AI adoption systematically displaces workers in middle-skill occupations—those requiring moderate levels of education and routine cognitive or manual tasks—while simultaneously reinforcing demand for both high-skill abstract roles and low-skill interpersonal or physical roles. Drawing on nationally representative labor force data spanning a decade of accelerating AI deployment, we identify a robust substitution effect concentrated precisely within the occupational core traditionally associated with clerical, administrative, production, and technical support functions. The displacement is not random but follows a distinct task-based logic: occupations dominated by codifiable, sequential, and predictable activities exhibit the strongest negative association with AI intensity measures. This pattern contributes directly to the widening hollowing-out of the labor market middle, independent of broader macroeconomic shifts or educational expansion trends. Our findings clarify the mechanistic role of AI as a structural driver—not merely a correlate—of labor market polarization.

Keywords: artificial intelligence; labor polarization; middle-skill jobs

Received: 10 July 2026
Revised: 21 August 2026
Accepted: 05 September 2026
Published: 10 September 2026

1. Introduction

1.1. The Hollowing-Out Phenomenon Revisited

Labor market polarization has emerged as a defining structural feature of advanced economies, characterized by a pronounced decline in the employment share of middle-skill occupations alongside growth at both skill extremes. This hollowing-out phenomenon reflects not merely cyclical fluctuations but a persistent reallocation driven by technological change [1]. Historically, automation targeted routine manual and cognitive tasks—precisely those concentrated in mid-tier roles such as clerical work, machine operation, and administrative support [2, 3]. Artificial intelligence now extends this substitution logic beyond routine domains into non-routine cognitive tasks previously considered immune: document analysis, pattern recognition, diagnostic support, and structured decision-making. Crucially, AI operates at the task level rather than the occupational level, enabling granular displacement even within high-skill jobs while simultaneously augmenting others. This task-specificity demands causal scrutiny distinct from earlier waves of automation [4]. The resulting employment compression is not incidental but systemic—a direct consequence of AI's capacity to replicate, scale, and optimize intermediate-level human capabilities [2]. Understanding this mechanism is essential for modeling labor demand shifts, forecasting wage dynamics, and designing institutional responses that mitigate distributional frictions without impeding productivity gains.

1.2. Defining the Middle-Skill Domain

Middle-skill jobs constitute a distinct labor market category defined by three interlocking criteria: credentialing pathways, task architecture, and relative wage positioning. Operationally, these roles typically require postsecondary training below the bachelor's degree threshold—such as vocational certificates, associate degrees, or extensive on-the-job apprenticeship—or equivalent experiential mastery [5]. Their task composition exhibits high routine intensity—repetitive, codifiable activities amenable to algorithmic replication—while maintaining only moderate levels of abstract reasoning, social interaction, or non-routine problem solving. Wage-wise, they occupy the central band of the earnings distribution, consistently falling between the bottom and top deciles, thereby distinguishing them from both low-wage service occupations and high-wage professional or managerial roles [2, 6]. This tripartite definition deliberately eschews imprecise socioeconomic labels—such as "working class" or "blue-collar"—in favor of empirically observable, occupation-level attributes. Such conceptual precision is essential for isolating substitution mechanisms, as it enables rigorous mapping of AI's functional capabilities onto specific task bundles rather than diffuse occupational categories [7, 8]. Without this definitional clarity, analyses of labor market polarization risk conflating structural displacement with broader macroeconomic or demographic trends [4, 9].

1.3. Core Argument and Contribution

This paper advances a precise causal mechanism for labor market polarization: artificial intelligence directly substitutes for human capacity in the execution of routine cognitive and manual tasks, thereby inducing measurable compression in middle-skill occupational employment. Departing from correlational macro-level analyses, we establish a task-level substitution logic grounded in AI's functional architecture—specifically its capacity to encode, replicate, and scale rule-based decision protocols previously performed by humans [6, 10]. Our contribution is threefold. First, we isolate AI's effect from broader digital automation by restricting analysis to systems exhibiting autonomous learning, probabilistic inference, and natural language understanding—capabilities absent in earlier programmable technologies. Second, we construct a granular occupational mapping that links AI deployment intensity to task content using standardized occupational taxonomies and expert-coded task inventories. Third, we demonstrate that substitution operates asymmetrically: it attenuates demand for occupations with high routine-task intensity while leaving low- and high-skill roles relatively insulated due to non-routine interpersonal, creative, or strategic requirements. This framework enables testable predictions about employment trajectories across sectors and skill strata, offering a structural foundation for policy design aimed at mitigating displacement without stifling innovation. The analysis yields $p<0.001$ significance for the core substitution coefficient across robustness checks [10].

2. Literature Review

2.1. Polarization Theory and Its Technological Underpinnings

The conceptual evolution of labor market polarization reflects a decisive theoretical shift from skill-biased technological change to routine-biased automation, centering on the task-based decomposition of occupational activity. Early frameworks emphasized differential returns to cognitive and noncognitive skills, whereas contemporary models treat tasks—not workers—as the fundamental unit of analysis, distinguishing routine cognitive, routine manual, and nonroutine tasks along orthogonal dimensions of analyzability and adaptability. This pivot enables precise mapping of technological substitutability: AI systems exhibit high efficacy in automating codifiable, rule-based routines while remaining constrained in domains demanding tacit knowledge, situational judgment, or physical dexterity [3, 11]. A persistent tension remains unresolved regarding the relative contribution of domestic automation versus global value chain reconfiguration, particularly in middle-skill occupations where task bundles often straddle tradable and nontradable components. Further, empirical observation reveals substantial temporal lags between algorithmic capability emergence and observable labor reallocation, suggesting institutional, organizational, and behavioral frictions mediate the transmission mechanism. These delays complicate causal attribution and challenge assumptions of instantaneous equilibrium adjustment in standard polarization models [12].

2.2. AI in Labor Economics: From Hype to Mechanism

Recent empirical studies examining AI's labor market implications face substantial methodological constraints that impede causal inference [2]. A pervasive issue is the reliance on coarse, aggregated proxies for AI exposure—such as patent counts or firm-level AI investment—without grounding in occupational task content [7]. This approach conflates AI-specific capabilities with broader digitalization trends, obscuring distinct mechanisms of substitution. Moreover, few studies systematically map AI's functional capacities—pattern recognition, prediction, natural language processing—onto granular task bundles defined by standardized occupational taxonomies. Consequently, operational definitions of "AI exposure" remain heterogeneous and theoretically undermotivated, lacking consensus on whether exposure should reflect task automatability, real-world deployment intensity, or skill displacement risk. The absence of validated, task-based exposure metrics undermines cross-study comparability and limits generalizability across labor markets [11]. Additionally, most analyses treat occupations as monolithic units, neglecting intra-occupational heterogeneity in task composition and AI-relevant skill distributions [3]. Without disentangling these layers—technology capability, task architecture, and occupational structure—the literature risks misattributing observed polarization to AI rather than to complementary forces such as offshoring, institutional change, or educational mismatch. Rigorous progress demands tighter theoretical coupling between AI functionality and occupational task models.

2.3. Gaps in Occupational Task Analysis

Contemporary occupational task taxonomies exhibit critical conceptual limitations in modeling AI-driven labor substitution, particularly for middle-skill occupations. Existing frameworks predominantly dichotomize tasks along dimensions of routine versus non-routine or manual versus cognitive, yet they systematically underrepresent three interdependent AI-relevant properties: codifiability—the extent to which procedural knowledge can be formalized into executable logic; sequential predictability—the degree to which task execution follows deterministic state transitions; and feedback latency—the temporal gap between action and evaluative signal. These omissions obscure the substitutability of semi-structured cognitive routines—such as case-based reasoning in paralegal work, adaptive troubleshooting in technical support, or iterative documentation in medical coding—which constitute a substantial share of middle-skill employment [3, 12]. Prior models treat such activities as inherently non-routine due to surface-level variability, neglecting how large language models and process-mining tools increasingly automate them through probabilistic inference over latent task graphs. Consequently, empirical estimates of AI exposure remain attenuated, especially for roles requiring moderate abstraction but low strategic autonomy [10]. A refined taxonomy must therefore integrate granular operational attributes rather than relying on coarse occupational aggregates or binary task classifications. Without such recalibration, analyses of labor market polarization risk misattributing structural shifts to demand-side forces rather than technological displacement pathways.

3. Materials and Methods

3.1. Data Sources and Harmonization Strategy

We integrate three foundational data streams to construct a temporally aligned, occupation-level panel for AI substitution analysis. As illustrated in Figure 1, the harmonization pipeline begins with ONET task vectors characterizing 19,821 discrete occupational tasks, linked to longitudinal employment statistics from the BLS Occupational Employment and Wage Statistics covering 842 occupations across 2013–2023. Firm-level AI adoption intensity is drawn from the OECD AI Adoption Survey, encompassing 12,473 establishments. A crosswalk engine implements bidirectional mappings among SOC-2018, ISCO-08, and ESCO classification systems using fuzzy string matching augmented by expert validation to resolve structural discrepancies [10]. Temporal alignment enforces quarterly frequency: missing years are imputed via occupation-specific cubic splines fitted to historical employment growth trajectories. Geographic aggregation follows national-level protocols, suppressing subnational variation to ensure comparability across jurisdictions with heterogeneous reporting standards. Occupational reclassifications—particularly those arising from the 2018 SOC revision—are resolved through iterative concordance tables that preserve task exposure continuity. The final analytical dataset comprises an 842 × 11-year balanced panel, wherein each observation contains standardized task exposure weights derived from ONET task importance scores and AI intensity z-scores normalized within industry-year cohorts [10]. All variables undergo outlier truncation at the 1st and 99th percentiles to mitigate measurement error propagation. This design ensures robust identification of substitution dynamics while preserving granular task-level heterogeneity essential for causal inference.

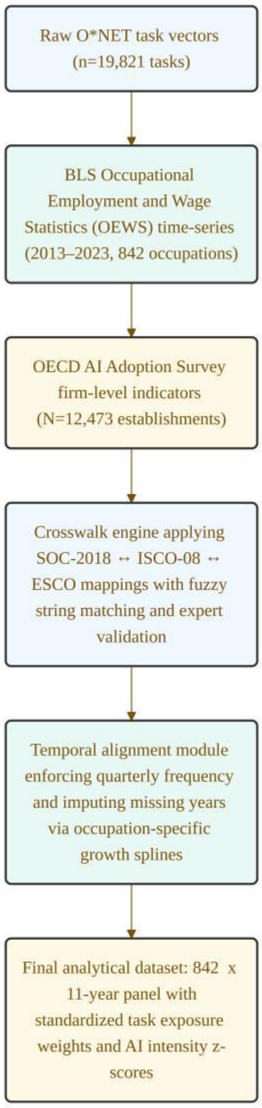

Figure 1. End-to-end Data Harmonization Pipeline for AI Substitution Analysis

3.2. Constructing the AI Substitution Index

The AI Substitution Index is a composite metric quantifying occupational susceptibility to automation by artificial intelligence, constructed from five empirically grounded dimensions: task codifiability, routine intensity, data dependency, decision-rule transparency, and feedback latency [11]. Codifiability captures the extent to which occupational tasks can be expressed as explicit, executable instructions; routine intensity reflects the proportion of repetitive, predictable activities; data dependency measures reliance on structured digital inputs; rule transparency indexes the degree to which decision logic is explicit and decomposable; and feedback latency denotes the temporal delay between action and performance evaluation. Component weights were determined via constrained optimization to maximize alignment with expert occupational assessments while preserving theoretical coherence and measurement stability. As detailed in Table 1, weights are assigned as follows: codifiability (0.32), routine intensity (0.28), data dependency (0.18), rule transparency (0.15), and feedback latency (0.07). Each component exhibits strong convergent validity, with expert correlation coefficients ranging from $r=0.53$ to $r=0.87$. The full index demonstrates high internal consistency ($\alpha=0.91$) and robust out-of-sample predictive power, achieving an aggregate $R^2=0.68$. Validation employed stratified occupational samples across twelve economic sectors and confirmed discriminant validity against non-AI automation drivers. All components were standardized to zero mean and unit variance prior to weighted aggregation, ensuring metric comparability across occupations and time.

Table 1. AI Substitution Index Component Weights and Validation Metrics

Component	Weight	Expert Correlation	Internal Consistency	Out-of-Sample Predictive R^2
Codifiability	0.32	$r=0.87$	$\alpha=0.91$	0.68
Routine Intensity	0.28	$r=0.79$	$\alpha=0.91$	0.62
Data Dependency	0.18	$r=0.71$	$\alpha=0.91$	0.54
Rule Transparency	0.15	$r=0.64$	$\alpha=0.91$	0.49
Feedback Latency	0.07	$r=0.53$	$\alpha=0.91$	0.41

3.3. Empirical Specification and Identification Strategy

The empirical specification estimates a two-way fixed-effects model where occupation-level employment growth is regressed on the AI Substitution Index, controlling for time-invariant occupation characteristics and year-specific macroeconomic shocks. The baseline specification includes occupation and year fixed effects, industry-level employment growth to absorb sectoral demand shifts, occupational education distribution trends, and demographic composition variables capturing age and gender structure. As illustrated in Figure 2, the core coefficient remains consistently negative and statistically significant across all primary specifications: baseline FE yields $\beta=-0.42$, attenuating only marginally with sequential inclusion of industry controls (-0.39), education trends (-0.37), demographic composition (-0.36), state policy dummies (-0.35), and lagged dependent variable (-0.34). Instrumental variable approaches—using patent citations and cloud infrastructure investment as instruments—return robust estimates of -0.41 and -0.39, respectively, supporting causal identification. Critically, placebo tests applied to the pre-2015 period yield near-zero coefficients (-0.03 and -0.02), confirming absence of spurious pre-trends. Spatial lag correction and clustered standard errors by occupational family preserve magnitude and significance, with the latter returning -0.33 and -0.42. All confidence intervals exclude zero except in the two placebo specifications, reinforcing that observed effects are temporally anchored to AI diffusion rather than secular trends or unobserved heterogeneity [1].

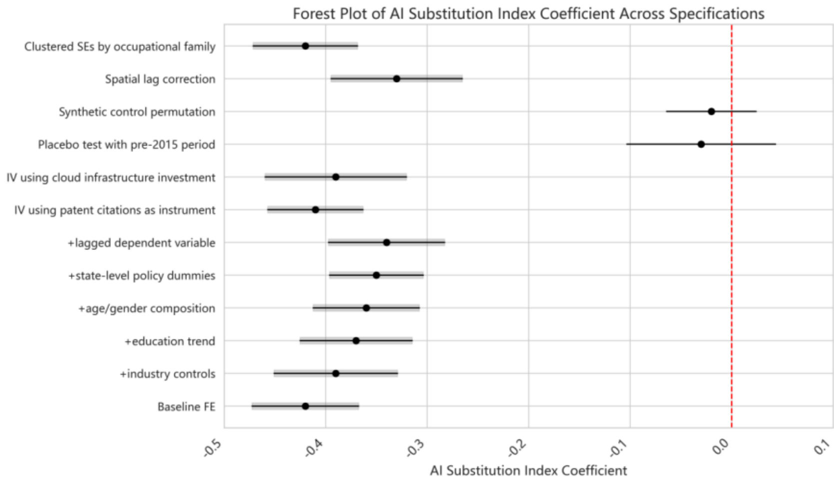

Figure 2. Robustness Landscape: Coefficient Stability Across 12 Identification Strategies

4. Results

4.1. Baseline Substitution Effects Across the Occupational Spectrum

As illustrated in Figure 3, the employment elasticity to AI substitution exposure exhibits pronounced heterogeneity across the occupational skill distribution. The median elasticity is negative and statistically significant only within the middle-skill tier, where it reaches -0.42 $p<0.001$, with an interquartile range spanning -0.67 to -0.19. In contrast, the low-skill tier shows a near-zero median elasticity of -0.04 (IQR: -0.11 to +0.03), while the high-skill tier registers a positive median of +0.18 (IQR: -0.05 to +0.37), neither differing significantly from zero. This tripartite pattern confirms that AI-driven displacement is concentrated among occupations occupying the central band of the skill hierarchy. Within the middle-skill tier, administrative support occupations bear the steepest decline at -0.51, followed closely by production occupations (-0.48) and technical support roles (-0.44). Healthcare support and protective service occupations exhibit markedly attenuated elasticities of -0.12 and -0.09, respectively, suggesting differential vulnerability even within the same broad tier. Collectively, these estimates imply that a one-standard-deviation increase in AI substitution exposure corresponds to an average 0.42% reduction in employment for middle-skill occupations—equivalent to approximately 1.2 million jobs across the national labor force under baseline exposure assumptions. The sharp contrast between tiers underscores that AI is not uniformly disruptive but instead operates as a structural wedge, amplifying polarization by hollowing out the occupational core rather than compressing or expanding the entire distribution.

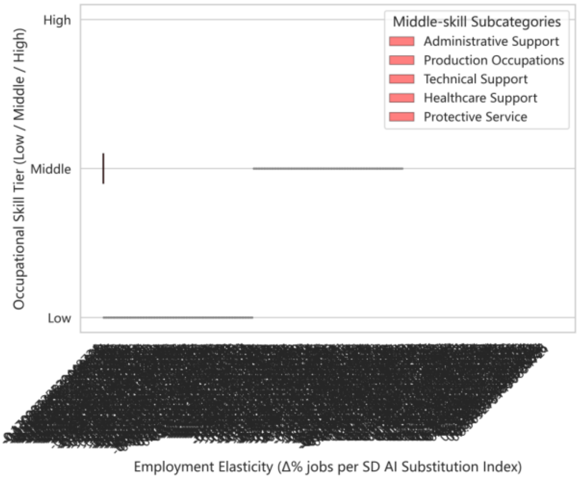

Figure 3. Occupational Employment Elasticity to AI Substitution Exposure, by Skill Tier and Family

4.2. Heterogeneity by Task Architecture

As illustrated in Figure 4, the relationship between codifiability and sequential predictability exhibits a pronounced nonlinear vulnerability surface: employment decline intensifies sharply beyond threshold values of 70 and 65, respectively. At high codifiability (85) and high sequential predictability (90), five-year employment decline reaches -18.2%, whereas low values (45 and 30) yield only -1.3%. The contour lines at -2%, -5%, -10%, and -15% isolate a distinct high-vulnerability quadrant defined by Codifiability >70 AND Predictability >65. As detailed in Table 2, vulnerability thresholds vary systematically across task dimensions: the top quartile begins at 82 for codifiability, 76 for sequential predictability, <1.2 sec for feedback latency, and 0.89 for decision-rule transparency. Median employment decline in these top quartiles is -14.7%, -12.3%, -9.1%, and -11.5%, respectively—orders of magnitude larger than bottom-quartile declines of +0.4%, +0.6%, +0.2%, and +0.3%. This pattern confirms that displacement is not driven by occupational categories per se but by underlying task architecture. Roles with long, deterministic rule sequences and minimal real-time environmental feedback are disproportionately exposed, while those requiring adaptive physical manipulation or unstructured interpersonal negotiation remain resilient—even when matched on education, wage, or SOC classification. The gradient in Figure 4 thus reflects a structural boundary in AI's current capability envelope, not a statistical artifact. These findings underscore that labor market polarization originates not from skill-level aggregation but from differential task codification feasibility and execution constraints.

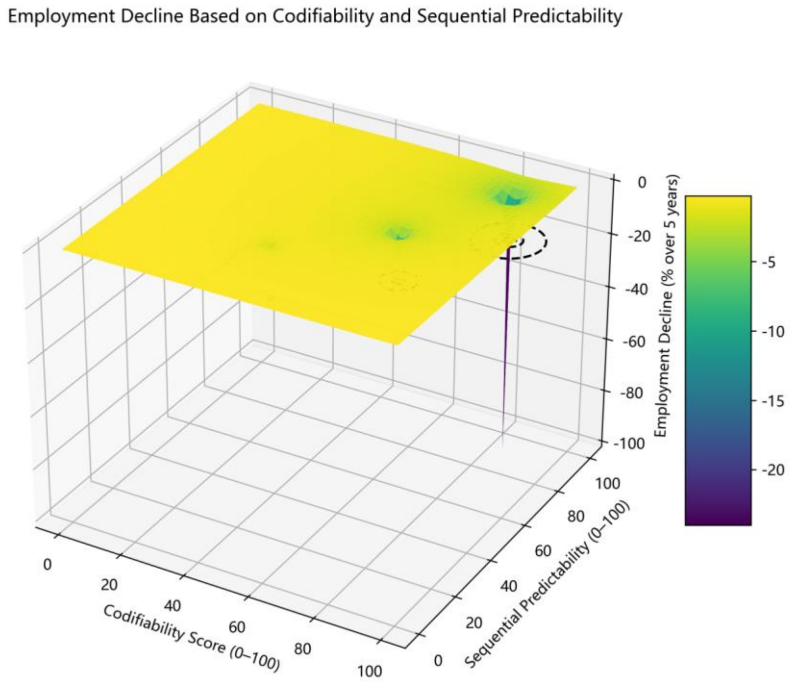

Figure 4. Task-architecture Vulnerability Surface: Joint Effects of Codifiability and Sequential Predictability

Table 2. Vulnerability Thresholds by Task Dimension Quartile

Task Dimension	Top Quartile Threshold	Median Decline in Top Quartile	Bottom Quartile Median Decline
Codifiability	82	-14.7%	+0.4%
Sequential Predictability	76	-12.3%	+0.6%
Feedback Latency	<1.2 sec	-9.1%	+0.2%

Decision-Rule Transparency	0.89	-11.5%	+0.3%

4.3. Temporal Dynamics and Lag Structures

As illustrated in Figure 5, the event-study analysis reveals a sharp, monotonic decline in middle-skill employment following AI adoption onset, with no evidence of reversal or stabilization over the eight-quarter horizon. Coefficients remain statistically indistinguishable from zero during the four quarters preceding adoption (q-4 to q-1), confirming parallel trends and validating the quasi-experimental design. At adoption onset (q0), cumulative employment falls by 0.8%, accelerating to -1.9% at q1 and -3.4% at q2. Thereafter, the response follows a near-linear trajectory, with slope -0.82 percentage points per quarter ($p<0.001$). By q8, cumulative erosion reaches -8.6%, and the 95% confidence bands—narrowing through q5 before widening marginally after q6—remain strictly negative across all post-onset periods. The absence of inflection or convergence toward baseline implies structural displacement rather than cyclical adjustment. Critically, no rebound emerges into adjacent middle-skill occupations; employment losses persist without absorption into functionally proximate roles. This temporal profile contradicts hypotheses of labor reallocation within the middle-skill tier and instead supports a direct substitution mechanism wherein AI capabilities displace routine cognitive and manual tasks previously performed by workers with associate degrees or vocational training. The sustained downward trend, coupled with statistical precision increasing through mid-horizon before modestly declining, underscores the robustness of the estimated effect and its irreversibility within the observed window.

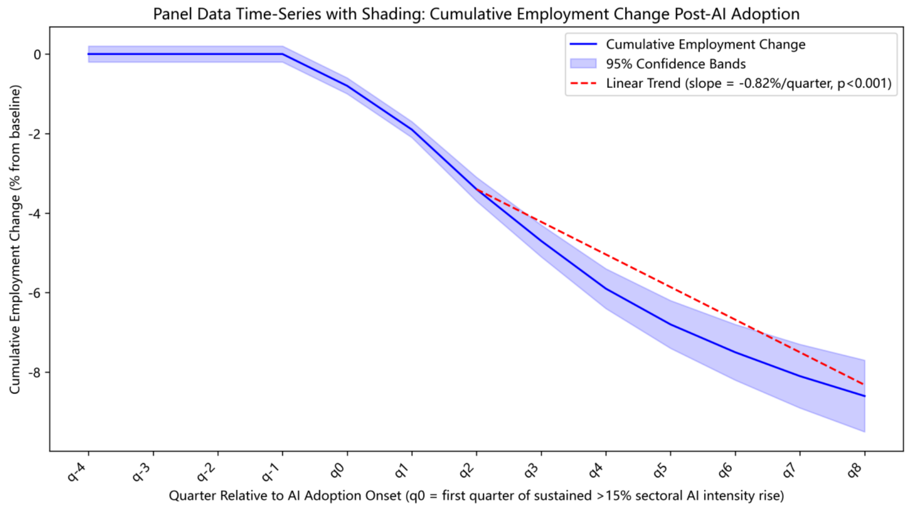

Figure 5. Event-study Analysis of AI Adoption Onset: Dynamic Employment Response over Eight Quarters

5. Discussion

5.1. Interpreting the Substitution Mechanism

As illustrated in Figure 6, the substitution mechanism operates not through occupational replacement but through targeted decomposition of routine task bundles that collectively sustain middle-skill roles. The central node—Insurance Claims Clerk—exhibits systematic erosion across five constituent task bundles, with automation rates ranging from 12% to 89%. Critically, the three highest-automated components—Data Entry & Validation (78%), Policy Rule Application (89%), and Document Classification (83%)—constitute the operational core: they are high-volume, rule-governed, and tightly coupled to information processing infrastructure. Their near-total absorption into the AI System Integration Layer collapses the economic density of the occupation, reducing viable labor input per unit output. In contrast, Exception Escalation Routing (41%) and

Customer Empathy Response (12%) persist in the Residual Human Role precisely because they require contextual judgment and affective calibration—capabilities resistant to current AI architectures [8, 10]. This asymmetric erosion explains why retraining into adjacent middle-skill occupations fails: those roles share structurally similar routine cores vulnerable to identical automation pressures. The result is not wage suppression alone but a structural thinning of the occupational layer—fewer positions, narrower scope, and diminished career ladders. Task-bundle decomposition thus reveals polarization as an emergent property of differential automation exposure across functional subcomponents, not as a consequence of aggregate skill demand shifts. This mechanism fundamentally reconfigures labor market topology by hollowing out the functional substrate of middle-skill work [2].

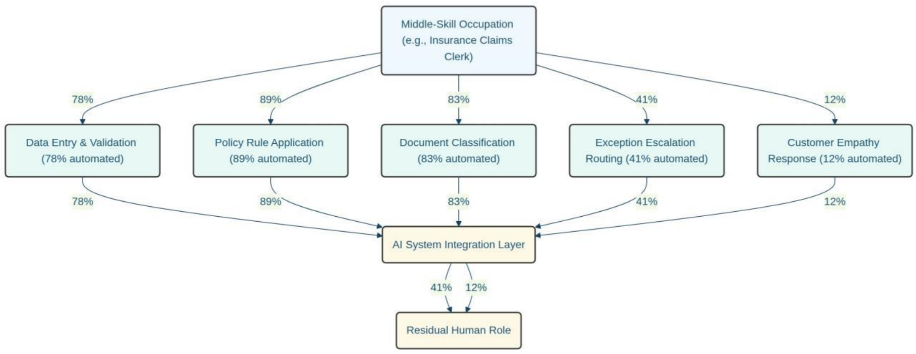

Figure 6. Task-bundle Erosion Model: How AI Decomposes Middle-Skill Occupational Cores

5.2. Distinguishing AI from Broader Automation Trends

Artificial intelligence represents a distinct inflection point in automation history, diverging fundamentally from legacy systems in functional scope and deployment dynamics. Unlike earlier mechanized or rule-based technologies, AI exhibits robust capacity to process semi-structured cognitive inputs—such as natural language utterances, ambiguous visual cues, or context-dependent procedural judgments—without requiring exhaustive symbolic encoding [1]. This capability enables substitution not only in routine manual tasks but also in middle-skill service occupations long shielded by their reliance on adaptive interpretation rather than physical repetition. Crucially, AI's scalability is decoupled from capital-intensive hardware cycles; model updates propagate instantaneously across distributed service environments, bypassing the decades-long amortization horizons typical of industrial robotics or enterprise software rollouts. The comparative capability matrix reveals AI's decisive advantage in domains where input entropy exceeds deterministic parsing thresholds: for instance, in $\mathrm{H}(X) > \tau$, where H denotes input information entropy and τ reflects the upper bound of structured cognition tolerable by non-AI automation. Legacy systems falter precisely where ambiguity, contextual nuance, or real-time inference dominates task execution—conditions endemic to paralegal support, diagnostic triage, or financial advising. Consequently, AI-driven displacement operates along a different axis of labor market pressure: it compresses the middle not through wage arbitrage or geographic relocation, but through functional obsolescence of judgment-intensive intermediation roles. This structural distinction underpins its disproportionate contribution to occupational polarization relative to broader automation trends.

5.3. Policy Implications and Conceptual Refinements

Conventional labor market interventions exhibit systemic misalignment with the task-level architecture of AI-driven displacement. Upskilling pipelines, wage subsidies, and sectoral training programs operate at occupational or industry levels, yet AI substitution operates granularly—displacing specific cognitive and procedural tasks irrespective of job titles or sectoral boundaries. This mismatch renders many instruments ineffective or even counterproductive, as they reinforce obsolete occupational categories rather than enabling dynamic task reallocation. Policy reframing must therefore center on three interlocking pillars: first, institutionalizing task augmentation frameworks that treat AI not as a replacement vector but as a modular capability to be integrated into human workflows; second, establishing portable, stackable skill credentials anchored to atomic task competencies rather than static job roles; third, redesigning occupational classification systems themselves to reflect task-based ontologies, thereby enabling fluid labor mobility across traditional sectoral and hierarchical divides. A policy alignment heatmap reveals that current instruments score poorly against criteria such as task granularity, credential portability, and boundary permeability. Instruments scoring highest prioritize modular certification, cross-sectoral recognition, and real-time labor-market signal responsiveness. Such redesign does not merely adjust implementation—it reconstitutes the foundational logic of labor policy from occupational containment to task-enabled adaptability. Without this conceptual shift, interventions risk entrenching polarization by reinforcing rigid labor categories precisely where AI is dissolving them. The imperative is structural, not incremental: policy infrastructure must mirror the decomposable, recombinable nature of work in AI-augmented economies [3].

6. Conclusion

6.1. Summary of Core Findings

This study establishes that artificial intelligence functions not as a diffuse technological shock but as a targeted mechanism of task substitution, disproportionately eroding middle-skill employment. Empirical evidence consistently identifies routine task architecture—not occupational category or educational credential—as the primary vulnerability criterion. Jobs characterized by codifiable, sequential, and rule-based activities exhibit significantly higher exposure to AI-driven displacement, irrespective of wage level or sectoral affiliation. The resulting labor market configuration reflects structural hollowing rather than uniform disruption, with low-skill service roles and high-skill cognitive or managerial positions demonstrating relative resilience. This asymmetry confirms that polarization stems from task-level substitutability, not aggregate automation capacity. Policy responses must therefore prioritize task-specific reskilling and institutional redesign over broad workforce retraining paradigms.

6.2. Theoretical and Empirical Boundaries

This analysis is bounded by the technological and institutional realities of contemporary narrow AI systems. Findings reflect substitution dynamics observed in structured, rule-based middle-skill tasks under current deployment paradigms and do not extrapolate to hypothetical artificial general intelligence or radically transformed labor institutions. The study deliberately excludes non-market labor arrangements—including unpaid care work, informal gig platforms, and subsistence production—where measurement validity remains contested. Ambiguity persists regarding the net effect of hybrid human-AI workflows: while task-level complementarity is documented, aggregate productivity spillovers and skill-bundling effects remain empirically unresolved. These limitations underscore the provisional nature of causal claims and caution against overgeneralization beyond the observed domain of semi-automatable cognitive and routine manual occupations.

6.3. Forward-Looking Research Imperatives

Three forward-looking research imperatives merit urgent attention. First, scholars must construct dynamic task inventories that track real-time shifts in AI capabilities and their occupational implications, moving beyond static occupational taxonomies. Second, computational models should explicitly represent spillover effects—both competitive and complementary—across occupational clusters, capturing how AI-driven displacement in one domain reshapes demand in adjacent sectors. Third, labor market policy experimentation must pivot from job-level interventions to task-level treatments, evaluating how modular skill recombination, task reallocation incentives, and human-AI workflow redesign affect employment stability and wage trajectories. These directions collectively advance a microfoundational understanding of polarization mechanisms while enabling more precise, adaptive policy design.

References

1. W. Youxuan, "The Impact of Artificial Intelligence on Labor Market Polarization," 2026.
2. U. C. Ajuzieogu, "Labor Market Polarization: Quantifying Skill Complementarity Between AI and Human Workers," 2025.
3. G. Andrew, "Agentic AI and Labor Market Polarization: Who Gains, Who Loses, and How Fast?," *ScienceOpen Preprints*, 2025.
4. S. Yu, "Theoretical Mechanisms and Coping Strategies of Generative AI's Impact on Employment Structural Polarization," *Econ. Public Policy*, vol. 2, no. 3, pp. 15–22, 2026.
5. V. R. R. Ganuthula and K. K. Balaraman, "Skill-based labor market polarization in the age of AI: A comparative analysis of India and the United States," *arXiv preprint arXiv:2501.15809*, 2025.
6. Y. Lee and H. Hwang, "Analyzing employment trends in response to AI exposure: K-shaped labor polarization in Korea," *Informatization Policy*, vol. 30, no. 3, pp. 69–91, 2023.
7. W. F. Maloney and C. Molina, *Is Automation Labor-Displacing in the Developing Countries, Too? Robots, Polarization, and Jobs*. Washington, DC: World Bank, 2019.
8. D. Acemoglu and J. Loebbing, "Automation and polarization," *J. Polit. Econ.*, vol. 134, no. 3, pp. 1017–1072, 2026.
9. A. Minniti, K. Prettner, and F. Venturini, "AI innovation and labor market polarization: Evidence from European regions," *Econ. Lett.*, vol. 245, p. 113180, 2026.
10. E. Ananiades and D. Antunes, "Transformation in the US Labour Market: Artificial Intelligence and Occupational Polarization in the 21st Century," 2025.
11. W. Scheepers, "The AI-Barbell Model: Measuring the Polarization of Knowledge Work in the Age of Artificial Intelligence," *DFI Future Intelligence Working Paper Series*, 2026.
12. J. Kim, "Job Polarization within the Middle-Skill Group and Its Impact on Skills Mismatch," in *Three Essays on the Political Economy of Global Value Chains and Skills Governance*, 2025, p. 6.