HARVARD Kennedy School
MOSSAVAR-RAHMANI CENTER
for Business and Government

Future of Work in the Age of Automation, Augmentation, and Agentic AI

How to Prepare the Workforce In a Post-AI Economy

Sol Rashidi
Harvard Kennedy School

June 2026

M-RCBG Associate Working Paper Series | No. 276

The views expressed in the M-RCBG Associate Working Paper Series are those of the author(s) and do not necessarily reflect those of the Mossavar-Rahmani Center for Business & Government or of Harvard University. The papers in this series have not undergone formal review and approval; they are presented to elicit feedback and to encourage debate on important public policy challenges. This paper is copyrighted by the author(s). It cannot be reproduced or reused without permission. Pursuant to M-RCBG's Open Access Policy, this paper is available to the public at hks.harvard.edu/centers/mrcbg free of charge. Papers may be downloaded for personal use only.

Mossavar-Rahmani Center for Business & Government
Weil Hall | Harvard Kennedy School | www.hks.harvard.edu/mrcbg

THE FUTURE OF WORK IN THE AGE OF AUTOMATION, AUGMENTATION, AND AGENTIC AI

How to Prepare the Workforce In a Post-AI Economy

Sol Rashidi, Senior Fellow

Harvard Kennedy School | M-RCBG Senior Fellowship

Research Assistants:
Anchita Batra, Sarang Agrawal, Amin Karimi Dehnashi, Kieran McCormick, and Johnny Nguyen

Working Draft

Last Updated: June 5, 2026

Table of Contents

Abstract 4

1.1 Introduction: Whether This Disruption Is Different 6

1.1.1 The Forecast Gap: Predicting the Economic Impact of AI 6

1.1.2 A Historical Perspective of Workforce Transformation 6

Seven Phases of Workforce Evolution 6

What the Seven Phases Reveal 8

1.2 The Scale and Pace of Displacement 8

1.2.1 Quantifying the Disruption: What the Data Actually Says 8

1.2.2 The Entry-Level: The Net Job Deficit Window 9

1.2.3 Impacts Across Geographies and Industries 11

Geographic Divergence: Uneven Exposure, Uneven Capacity to Respond 12

Industry-Level Divergence: Where AI Disrupts First, Fastest, and Most 14

The Common Trend 16

1.3 Intellectual Atrophy and the Competence Crisis 16

1.3.1 How Humans Build Expertise And Why AI Is Disrupting this Pathway 16

1.3.2 Cognitive Offloading: The Neurological and Psychological Evidence 17

1.4 Conclusion: From Displacement to Developmental Risk 20

2.1 The Case for Immediate Intervention 22

2.1.1 Reimagining Human Intelligence in an AI-Augmented World 22

2.2 Blueprint for Deliberate Redesign for Businesses 23

2.2.1 The New Talent Taxonomy 23

The M-Shaped Talent Taxonomy Model 23

2.2.2 Role Archetype Considerations for the AI Era 27

Why a Typology 27

Real-World Example: Search Marketing Strategist 32

Boundary Conditions 34

2.3 Legislative Challenges, Policy Change and Organizational Imperatives 34

2.3.1 Can Education and Reskilling Help? 34

2.3.2 Why Markets Will Not Self-Correct By Themselves 35

2.3.3 The Legislative Challenge: Speed, Complexity, and Competing Interests 36

2.3.4 Eight (8) Frameworks and Policy Changes 36

The AI Nutrition Label: Framework for Data Transparency 37

AI Agent Payroll Tax 37

Early-Career Hiring Tax Credits: Developing the Experts of the Future 38

AI Deployment Registries 39

AI-ISO Standards 40

Extend Antitrust Laws to Private AI companies 41

Displacement Notification Requirements....41
Autonomous Capability Threshold (Kill Switch)....42
Policy Change Summary....42

2.4 Preserving Human Capability in the AI Economy....44

2.4.1 What the Evidence Demands....44

2.4.2 The Workforce 4.0 Imperative....45

References....47

Appendix A: Methodological Breakdown of the Post-AI Pay Structure....65

Appendix B: Methodological Walkthrough — Task Scoring and Sensitivity Analysis....72

B.1 The Eloundou exposure rubric....72

B.2 Published scores for search marketing and related occupations....73

B.3 Derivation of the three productivity options....74

B.4 Caveats and boundary conditions....75
Three methodological caveats apply.....75

B.5 The 36 tasks of Search Marketing Strategist: full assignment to Eloundou category and OOA tier....75

Appendix C: The OOA Framework Applied, Two Contrasting Functions....82
Function 1: Clinical Decision Support — An Artisan-Anchored Function....82
Function 2: Credit Underwriting — An Orchestrator-Anchored Function....84
Synthesis: What the Two Examples Demonstrate....86

Appendix D: This Paper's Framework Among Other Frameworks....87

D.1 The two-prescription gap matrix....87

D.2 The three practitioner frameworks engaged in depth....92
McKinsey Agentic Organization....92
Deloitte 2025 Global Human Capital Trends and Humans × Machines....93
Daugherty & Wilson Human + Machine....94

D.3 A side-by-side worked example: Search Marketing Strategists....95
What this paper's framework (OOA + M-shape) prescribes....95
What McKinsey Agentic Organization would prescribe....95
What Deloitte Humans × Machines would prescribe....96
What Daugherty & Wilson Human + Machine would prescribe....96

Abstract

Every major technological disruption in history automated routine physical or procedural tasks. Artificial intelligence is the first technology capable of automating meaningful portions of cognitive work, particularly the entry-level and developmental work through which expertise, judgment, and professional competence have traditionally been formed. As a result, AI represents more than another wave of productivity-enhancing technology. Through automation, augmentation, and autonomous agents, it has the potential to reshape the mechanisms by which organizations develop talent, transfer knowledge, and cultivate future experts. The central concern of this thesis is therefore not workforce displacement alone, but the preservation of the capability-formation systems that sustain expertise, human judgment, and long-term workforce resilience in increasingly complex and uncertain environments.

Prior disruptions displaced workers from tasks while preserving the experiential pathways through which new forms of expertise emerged. New jobs were created. Labor markets adjusted. The developmental ladder held. AI disrupts that sequence at its foundation. By automating the entry-level and early-career work that has historically served as the training ground for professional judgment, AI is not merely eliminate roles, it's dismantling the pipeline through which the next generation of experts, supervisors, and decision-makers are produced, which are needed to oversee and govern the AI and business ecosystems of the future. While the consequences are not yet visible in aggregate employment statistics, they are visible in the numbers that precede them: a 6–13% employment decline among workers aged 22–25, a 15–20% drop in graduate-level job postings, and a 14% decline in the monthly rate at which young workers successfully enter those occupations. We establish that a *talent formation fracture* is underway and this paper responds to that fracture on three (3) fronts.

It documents the scale and cost of AI-driven disruption, including the pathway through which cognitive capability erodes when the work that builds it disappears. It explores the velocity of AI-driven labor transformation, sectors and geographies most exposed, as well as the mechanisms through which entry-level and experiential work are being compressed. It also presents a unique analysis of cognitive and developmental consequences associated with the erosion of experiential learning, including the emergence of Intellectual Atrophy™ as a compounding risk factor. To address these challenges, this thesis introduces a set of workforce frameworks designed for Workforce 4.0—the workforce operating alongside automation, augmentation, and agentic AI. These frameworks provide a structural model for redesigning organizations in ways that preserve human capability while leveraging the benefits of AI. Central to this approach are the M-Shaped Talent Framework and the Orchestrator-Operator-Artisan (OOA) role architecture, both designed to strengthen the uniquely human capabilities that remain essential in increasingly AI-enabled environments. Together, they offer practical mechanisms for sustaining expertise development, workforce adaptability, and long-term organizational resilience.

The thesis further proposes seven (7) policy interventions intended to align the incentives of AI adoption with the long-term development of human capital and the future of work. Collectively, these recommendations seek to mitigate emerging risks across labor markets, education systems, and public institutions while ensuring that technological progress and workforce sustainability advance in parallel rather than in opposition. The opportunity for proactive action remains open, but it is unlikely to remain so indefinitely. The workforce patterns, organizational structures, and economic incentives emerging today will shape the capability landscape of future generations. Once embedded, these dynamics may become increasingly difficult to reverse. The objective of this thesis is therefore not to resist technological progress, but to provide a framework for managing it responsibly—preserving the systems through which expertise, judgment, and human capability are developed before they are permanently altered by the very technologies they helped create.

PART 1:
THE TALENT FORMATION FRACTURE, WORKFORCE OF THE FUTURE AND WHAT HUMAN CAPITAL MEANS IN THE AGE OF AI

1.1 Introduction: Whether This Disruption Is Different

1.1.1 The Forecast Gap: Predicting the Economic Impact of AI

The task of forecasting AI's economic impact has produced an unusually wide variance in expectations; fundamental disagreements about the relationship between AI capability and human labor, and there's no reconciling the discrepancy because of capital-market narratives, and the absence of historical precedent for a general-purpose technology that advances at this velocity. This thesis characterizes that divergence as the *Great Forecast Gap*.

On one camp, leading AI developers from Anthropic Massenkoff and McCrory (2026), caution that prior exposure-based forecasts identify roughly a quarter of US jobs as vulnerable and suggest that the methodology used to measure AI risk may be underestimating the impact and by itself, could be a source of forecast error (Massenkoff & McCrory, 2026). On the other side of the spectrum lies a strongly optimistic view from Goldman Sachs Global Investment Research, estimating that widespread AI adoption could increase global GDP by approximately 7% over a decade, nearly $7 trillion in economic value, driven by productivity gains across both advanced and emerging economies.
MIT economist Daron Acemoglu reaches a starkly different conclusion. His 2024 NBER working paper projects no more than a 0.66% increase in total productivity over a ten-year horizon, arguing that a substantial portion of the modern knowledge economy work requiring context, tacit knowledge, adaptability, and social judgment, remains resistant to automation precisely because it lacks the codifiable outcome measures from which AI models reliably learn (Acemoglu, 2024). Additionally, the University of Chicago Booth School survey of U.S. economic experts also reflects this division empirically stating only 4% of respondence expressed strong confidence in AI-driven growth, 42% agreed, 47% remained uncertain or divided, and 8% disagreed (IGM Forum, 2023).

Accordingly, the central focus of this thesis is not the extent to which artificial intelligence is overhyped, nor the precise number of jobs that may be displaced through automation. Rather, it is the design of workforce, organizational, and institutional systems that can preserve human capital, and the future of work, upon which expertise, judgment, and economic resilience depend. As AI becomes increasingly autonomous, augmentative, and agentic, the challenge is not simply adapting work to intelligent machines, but ensuring that the systems responsible for developing human capability remain intact.

1.1.2 A Historical Perspective of Workforce Transformation

To assess whether AI constitutes a structural break or a continuation of established patterns, we examine the seven phases of workforce evolution up until 2022, each defined not only by the nature of the innovation but by the corresponding shifts in labor structure, skill formation, and institutional response.

Seven Phases of Workforce Evolution

Phase 1: The Invention of the Job (Late 1700s–1800s). Industrialization broke the agrarian model, pulling workers from fields into factories. Muscle power was traded for steam and machine power and the transition displaced agrarian work but generated net employment growth in factories needing more bodies, not fewer.

Phase 2: The Assembly Line (Late 1800s–Early 1900s). Electricity and standardized parts enabled mass production, simplifying work into discrete learnable steps. Craft workers became mass producers and labor demand expanded alongside productivity because simplification allowed enterprises to scale.

Phase 3: The Rise of the White-Collar Worker (1930s–1960s). Industrial complexity created demand for clerical oversight and management, creating a new layer of work. For the first time, education and credentials determined career trajectory. The workforce pyramid, wide at the base and narrow at the apex became the standard model, with each rung of the career ladder building on the competencies developed at the rung below. Hence, career progression through the *corporate ladder* was born.

Phase 4: The Digital Desk (1970s–1990s). The personal computer replaced clerical roles but expanded the overall workforce by 51%, with growth concentrated in services and professional functions. While technology automated and displaced specific tasks, it created new categories of work in parallel, and faster than it eliminated them.

Phase 5: The Digital Global Economy (1990s–2010s). The internet enabled offshoring and global services economy. Management roles expanded, mid-tier roles expanded, even though production roles declined. The concept of job security eroded as many things were digitized, but aggregate employment continued to grow.

Phase 6: Digital Transformation (2010s–Early 2020s). The boundary between technology jobs and non-technology jobs narrowed and every function was required to adopt digital fluency. Marketing merged with digital and finance merged with algorithmic modeling. Despite continuous disruption, employment rose across sectors.

Figure 1: Job evolution by industry group (million)

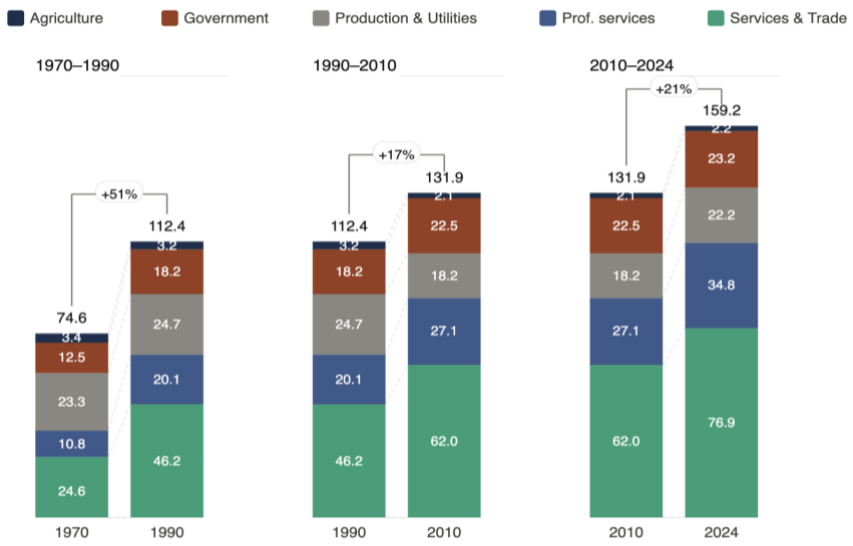

Note: This figure illustrates changes in employment across major macro-sectors during phases 4, 5, and 6 of workforce evolution. In most sectors, total employment either grew or remained relatively stable, with only a small number experiencing absolute declines. The figure also highlights the broader structural transition of the economy away from production-oriented work and toward a labor market increasingly concentrated in service-based industries (Bureau of Labor Statistics, n.d.-b, Federal Reserve Bank of St. Louis, n.d.-a, Federal Reserve Bank of St. Louis, n.d.-b).

Phase 7: The AI Revolution (Present). For the first time, automation targets how workers synthesize data into information and actional insights that drive business growth, rather than what they physically produce or process and intelligence is now becoming a scalable utility. Labor demand is declining across nearly all industries and of the jobs created in the midst, AI fluency and expertise is a pre-requisite, and entry-level and early-career adults are experiencing unemployment at an all-time high.

What the Seven Phases Reveal

The pattern across Phases 1 through 6 is consistent: technology automated specific forms of labor, displaced certain roles, and disrupted certain workers, but net employment expanded as new categories of work emerged. Crucially, each prior disruption left the developmental ladder intact. Even as specific tasks disappeared, new forms of tasks and work continued to provide the experiential grounding required for skill accumulation, judgment formation, and professional advancement. Workers moved up the pyramid, because the base of the pyramid kept regenerating.

AI introduces a different dynamic. Rather than targeting manual work, AI systems automate core cognitive tasks, analyzing datasets, synthesizing information, drafting documents, generating code, executing decision-support workflows, which is the foundational layer of white-collar, entry-level and early-career work. This creates a structural discontinuity absent in prior technological shifts. The critical question with this finding is therefore, not how much displacement will occur, or where displaced workers are reallocated, but how future workers will acquire the competencies required to oversee, interpret, and govern the AI systems of the future, particularly in environments that continue to rely on "human- experts-in-the-loop" as a core governance mechanism.

Under this new paradigm, the implications extend beyond short-term employment effects to the long-term threat of the value of human capital, expertise development pathways, and compression of the workforce pyramid. What emerges is not simply a redistribution of labor, but the risk of an underdeveloped talent pipeline, a deficit of expertise, and a workforce insufficiently prepared to work alongside emerging systems and one in which the value of human capital gets questioned. However, if we plan proactively, which means both governance protocols and policy changes enable a health transition for the workforce, we can develop a symbiotic relationship between humans and AI, as discussed by Soroush Saghafian in his paper, The Human-Algorithm Centaur (Saghafian, 2024), where there is a complimentary value that gets created when the power of both humans and AI joins forces.

1.2 The Scale and Pace of Displacement

1.2.1 Quantifying the Disruption: What the Data Actually Says

The scale of AI-driven workforce disruption is vast and broad. The World Economic Forum projects that AI and automation will displace approximately 85 million jobs globally while creating 97 million new roles, a net positive. However, the data obfuscates whether the workers displaced can and will have access to the jobs being created or if the net new jobs are technically oriented.

The International Monetary Fund (IMF, 2026) now estimates that nearly 40% of jobs globally are exposed to AI, with exposure rising to 60% in advanced economies, where cognitive and white-collar work is more prevalent. The World Economic Forum's 2025 Future of Jobs analysis indicates that tens of millions of roles will be displaced this decade, while a comparable number of new roles emerge, resulting not in simple net loss, but in significant reshaping of job composition and skill demand.

What differentiates the present moment from prior technological disruptions, industrialization, electrification, digitization — is not the existence of job loss, but its distribution of impact. Modern AI systems are uniquely capable of automating cognitive and linguistic tasks that sit at the core of white-collar and professional work, and these are disproportionately concentrated in entry-level and early-career roles. Data from Anthropic's own usage logs corroborate this pattern Massenkoff and McCrory (2026) found that computer programmers carry 75% task coverage, followed by customer service representatives

and financial analysts. Critically, they suggest displacement risks documented has not yet reached its ceiling.

1.2.2 The Entry-Level: The Net Job Deficit Window

The entry-level work has been referred to as the 'canary in the coal mine' of AI job displacement and one of the first areas to be impacted by automation (Stanford Digital Economy Lab, 2025). Traditionally, the tasks assigned to early-career adults are structured enough to delegate, straightforward enough to supervise, and frequent enough whereby daily practice creates experience which organizations require for performance at scale.

Empirical labor market data increasingly confirms this alignment. Studies of employment trends between late 2022 and mid-2025, when the adoption of AI rose, show that early-career work experienced employment declines of 6–13%, while employment for mid-career and senior workers in the same occupations remained stable or increased (Stanford Digital Economy Lab, 2025). Researchers analyzing payroll data report that workers aged 22–25 saw employment fall approximately 6% between late 2022 and September 2025, while older workers in those same occupations experienced employment growth of roughly 6–9% over the same period (Stanford Digital Economy Lab, 2025). Firm-level event studies show similar dynamics, with young workers experiencing a 16% relative employment decline (Stanford Digital Economy Lab, 2026). Unemployment among recent college graduates reached 5.6%, an all-time high (NY FED 2025). At the same time, graduate-level job postings declined by 15–20%, even though applications per posting rose sharply, indicating simultaneous contraction in labor demand and increased competition among job seekers (Handshake, 2025). A separate analysis by Anthropic, using Current Population Survey data reaches a comparable conclusion. Massenkoff and McCrory (2026) discover job-finding rates declined by approximately 14% in the post-ChatGPT period compared to 2022 baselines. Young labor market entrants may be leaving the labor force altogether rather than registering as unemployed, causing unemployment statistics to understate the magnitude of early-career disruption. In software development, entry-level hiring declined by roughly 20% even as demand for senior engineers increased. In consulting, accounting, and legal services, entry-level hiring has fallen between 15% and 25% since 2022 with similar dynamics are emerging in marketing analytics, corporate finance, and customer operations.

The common trend is AI is absorbing the execution-heavy tasks historically assigned to early-career workers, and organizations are continuing to respond by reducing early-career hiring (IMF, 2026). Employers themselves anticipate automation rates of roughly 65% for information and data processing tasks compared to 35% for reasoning and decision-making tasks, a ratio that maps closely onto the distinction between entry-level and senior-level work. Early-career adults in this sense, are not simply the first casualties of AI-driven displacement, they are the earliest and most visible signal of how AI-driven workforce restructuring is unfold broadly.

A second theme that is also resurfacing in the growth of new roles is they are concentrated in AI-adjacent and AI-enabling functions such as machine learning engineering, cybersecurity, and model risk management, all requiring technical skills and a certain degree of experience. In fact, the World Economic Forum projects that the fastest-growing job categories through 2030 include AI and machine learning specialists, and fintech engineers, a concentration in technical roles (WEF, 2025); and the data shows none of these roles are entry-level or early-career positions. They require advanced technical expertise and years of accumulated domain knowledge, precisely the kind of expertise that is built through the entry- level and early-career roles that are now being eliminated.

This thesis refers to this gap between jobs eliminated and jobs created as 'the net job deficit window' — a transitional period during which labor markets experience heightened friction, underemployment, and

polarization, even if long-term equilibrium can re-emerge. We believe the span of time to reach positive net impacts, if ever reached, will be greater and longer than prior deficit windows of past disruptions as indicated in the decline of graduate hires, entry-level posting volumes, and the diverging employment trajectories of workers in AI-exposed occupations.

Additionally, we estimate, the net job deficit window will create a deeper structural risk that the economy won't be able to rebound from, due to what this paper terms as "the talent formation fracture." This fracture emerges as AI absorbs work historically assigned to graduates and early-career professionals, effectively hollowing out the base of the workforce pyramid. As entry-level opportunities disappear, the pyramid compresses upward, weakening the talent-development pathways that organizations rely upon to cultivate future expertise. Catalini, Hui, and Wu describe this phenomenon as the "Missing Junior Loop," in which apprenticeship pathways collapse as AI absorbs entry-level work, shrinking the future stock of human expertise precisely when human oversight becomes more valuable as AI scales (Catalini et at., 2026). If "the talent formation fracture" is sustained across multiple graduating classes, as the current trajectory suggests, the consequences extend beyond individual careers and into organizational capacity to hire experts because the supply of experts will compress as demands will increase or remain constant.

We refer to this gap as 'the oversight gap', which means, in time, we will lack the experts needed to oversee and govern the orchestration of AI and business ecosystems in the next decade to come. Organizations that eliminate entry-level roles to capture short-term productivity gains may find, within five to eleven years, the time period in which expertise and mastery is developed, their pipeline of experienced talent has narrowed and its because they collectively eliminated the conditions under which mastery and expertise are produced.

Figure 2: The workforce pyramid and its potential evolution

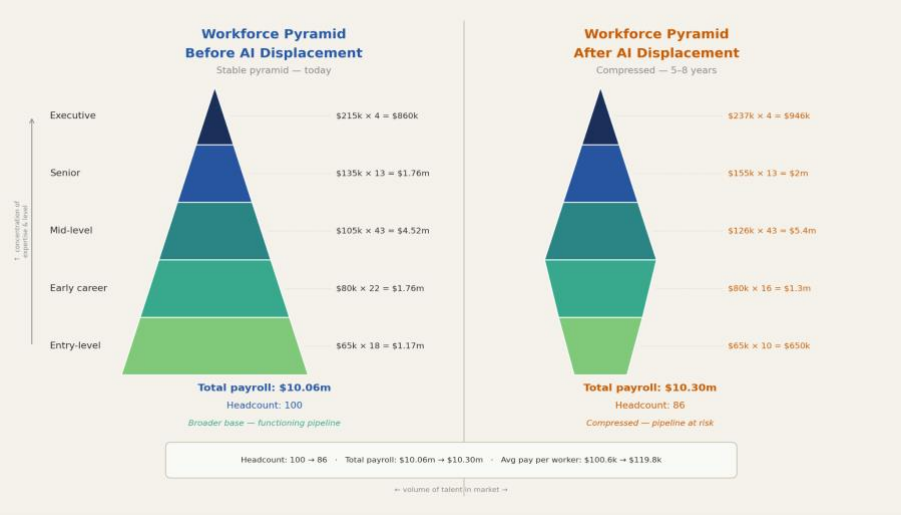

Note: This figure illustrates how AI-driven workforce compression may disproportionately reduce entry-level and early-career roles while increasing reliance on a smaller pool of experienced workers. As the pipeline for developing future talent narrows, expertise becomes more valuable and scarce, placing a growing premium on workers with advanced judgment, institutional knowledge, and oversight capabilities (Brynjolfsson et al., 2025; Dell'Acqua et al., 2023; Federal Reserve Bank of New York, 2026).

Figure 2 operationalizes the workforce-compression dynamic using a calibrated model of a representative 100-person white-collar, knowledge-work firm. The model begins with an approximated baseline split of roles, 18 percent entry-level, 22 percent early-career, 43 percent mid-level, 13 percent senior, and 4 percent executive, then applies AI-induced compression only to the bottom two layers, reducing entry-level headcount to 12 and early-career headcount to 17 while holding the upper layers constant. This

structure is consistent with the evidence reviewed above: AI is first absorbing execution-heavy junior tasks, while firms still retain experienced workers for judgment, accountability, and oversight (Brynjolfsson et al., 2025; Dell'Acqua et al., 2023; Federal Reserve Bank of New York, 2026).

Using rounded wage anchors of $65K, $80K, $105K, $135K, and $215K across the five layers, baseline payroll totals $10.06 million. In the compressed case, the model applies scarcity and complementarity premia of 15 percent for mid-level workers, 10 percent for senior workers, and 5 percent for executives. These premia reflect the logic that once many firms thin the same junior pipeline, the relative supply of experienced labor tightens and the cost of preserving output shifts upward toward workers with tacit knowledge and managerial judgment (Bidwell, 2011; Katz & Murphy, 1992; Acemoglu & Restrepo, 2019; Davis, 2026).

The result is structurally significant. Total headcount falls from 100 to 89, yet total payroll rises slightly to $10.16575 million, marginally exceeding baseline. Average payroll per worker rises from $100,600 to $114,222. The implication is not that every firm will immediately face a higher wage bill after reducing junior hiring. It is that AI-induced compression at the base of the pyramid should not be assumed to be a durable net cost reducer because firms still need to preserve their operations with more experienced, and more expensive workforce.

The deeper significance is, the choices made today that reduce short-term payroll savings, is the same process by which organization may experience higher payroll costs in five (5) to eleven (11) years. Figure 2 supports the paper's broader claim: the central risk is not only displacement in the short run, but the gradual erosion of the talent pipeline through which affordable, experienced human judgment is produced. Full model construction, assumptions, and parameter rationale are provided in Appendix A.

1.2.3 Impacts Across Geographies and Industries

The analysis thus far has focused primarily on evidence from the United States, where the earliest and most visible signs of AI-driven workforce restructuring is emerging. However, these dynamics will not unfold uniformly. While the underlying mechanisms of AI-driven disruption are broadly consistent, their impact will vary across geographies and industries depending on differences in economic structure, demographic composition, institutional capacity, and technological infrastructure. Geographic variation shapes the distribution of exposure and adaptive capacity across countries, while industry structure determines how quickly and deeply AI reshapes work within those economies.

Geographic Divergence: Uneven Exposure, Uneven Capacity to Respond

Figure 3: Geographical exposure matrix

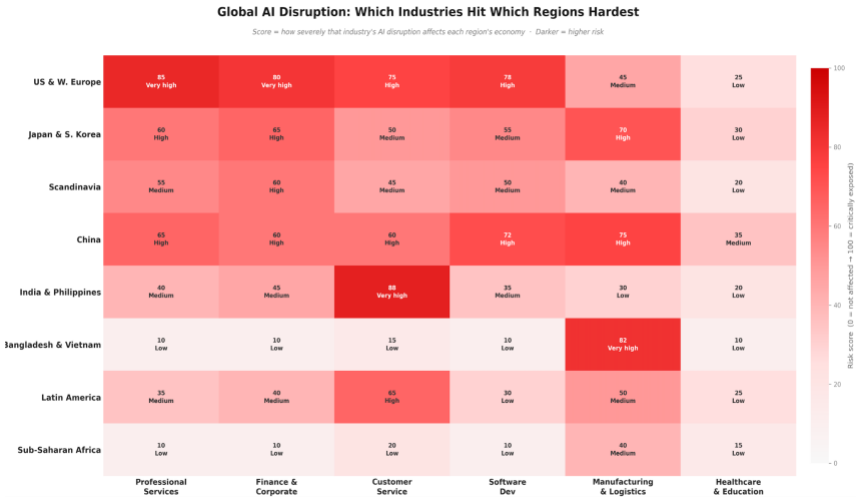

Note: This figure illustrates the uneven geographic exposure of labor markets to AI-driven disruption. Advanced economies face higher immediate exposure due to their concentration in knowledge-intensive and cognitive occupations, while emerging economies face greater long-term structural risk as AI threatens traditional industrialization and offshoring pathways before these economies have fully developed them (Garter, 2025; IEA, 2025; Gen-AI: Artificial Intelligence and the Future of Work, 2024; Klarna, 2025; Revelio Labs, 2025; Stanford HAI, AI Index, 2026; Future of Jobs Report, 2025).

To understand the distributional effects of AI, we have to consider how disruption varies geographically. While AI is a global technology, its labor market consequences are not uniform. The mechanisms described earlier, workforce pyramid compression, career-ladder thinning, and the talent formation fracture, operate across borders, but their impact depends on economic structure, demographic composition, institutional capacity, and technological infrastructure within the region. The result is not a single global labor shock but a set of regionally distinct disruptions unfolding at different speeds and with different capacities for adaptation.

Advanced Economies: Higher Exposure, But Higher Adaptive Capacity

Advanced economies, including the United States, Western Europe, Japan, and South Korea, face the highest immediate exposure to AI-driven labor disruption. Employment in these economies is concentrated in knowledge-intensive professional roles, the very domains where generative AI demonstrates the greatest capability. IMF estimates suggest that 40–60% of jobs in advanced economies are exposed to AI displacement (IMF 2026).

Emerging Economies: Lower Exposure, But Higher Structural Risk

Emerging economies face the opposite challenge: lower immediate exposure to AI but higher long-term structural risk as traditional development pathways, from agriculture to manufacturing, are disrupted before many countries have fully traversed them. AI-driven automation threatens two pillars of industrialization in developing economies. In manufacturing, robotics and AI will reduce the advantage of off-shore and near-shore low-cost labor, increasing the possibility of reshoring production to advanced economies creating massive economic downfall in developing and emerging markets. In services, AI-powered customer support and administrative automation threaten the offshoring sectors that have driven employment growth in countries such as India and the Philippines. Economists increasingly describe this

dynamic as 'premature deindustrialization.' Countries lose labor-intensive job market before capturing the productivity gains that industrialization historically delivered (Rodrik, 2016). These risks are compounded by their limited capacity to deploy the necessary programs to keep pace with the velocity of AI development.

Infrastructure constraints further widen the gap. AI development currently remains concentrated in advanced economies. In 2023, the United States attracted nearly nine times more private AI investment than China and vastly more than any developing economy (Stanford HAI, 2024). High-income countries host roughly 77% of global data center capacity, while low-income countries hold less than 0.1% (IEA, 2025). Without access to the energy, connectivity, and compute infrastructure required for AI deployment, many emerging economies risk absorbing displacement effects without capturing the productivity gains, exacerbating the phenomenon of the digital divide — the gap between individuals, households, businesses, geographic areas and their ability to access devices, high-speed connectivity, and digital literacy skills because of their socio-economic conditions.

In addition to geographic variances, demographic trends further intensify these differences. Advanced economies, such as Japan, face aging populations and shrinking workforces, AI adoption partially offsets labor shortages. In such contexts, a smaller entry-level workforce may not represent a crisis but a demographic adjustment.

Emerging economies however face the opposite problem. Regions such as Sub-Saharan Africa and parts of South Asia are in need of generating millions of new jobs annually to absorb rapidly growing youth populations. But with AI-driven efficiency gains in manufacturing, services, and administrative work, it further reduces the entry-level opportunities precisely when labor market demand is highest. In such contexts, the 'talent formation fracture' could extend beyond individual careers to the structural exclusion of entire demographic groups from the formal economy.

Figure 4: Evolution of population and job growth in productivity revolution leaders, *million people/jobs*

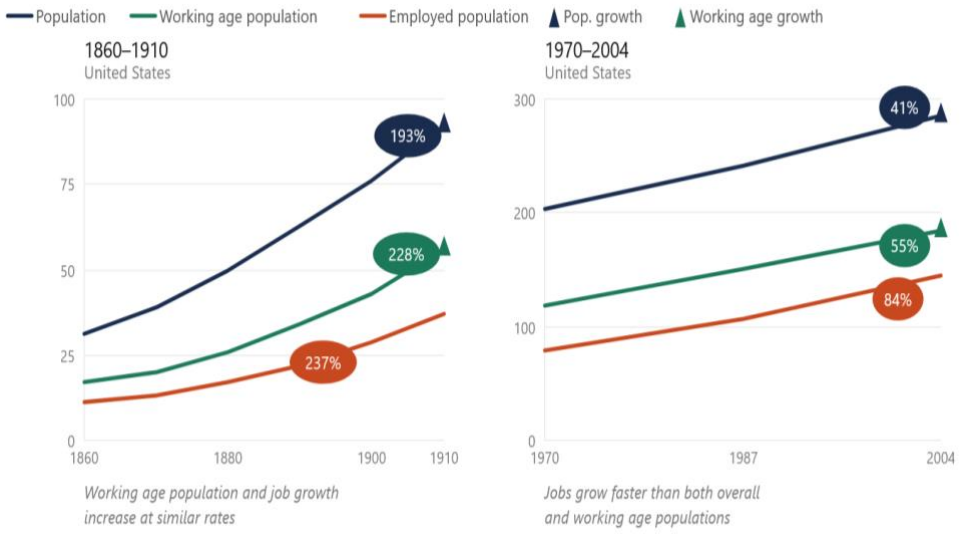

Note: This figure illustrates the trend growth in population and employment statistics across major productivity revolutions (Voth, 2001;Lebergott, 1966;Gibson, 2010;United Nations DESA, 2024;U.S. Department of Labor, n.d.).

Industry-Level Divergence: Where AI Disrupts First, Fastest, and Most

While geographic and demographic differences shape the distribution of AI's impact, industry-level differences further determine the speed and depth of disruption within those economies. Sector variations in task composition, data analysis, synthesis, and repetitive process flows determine whether AI acts primarily as an augmentation tool or as a mechanism for labor substitution.

Professional Services

Professional services, including consulting, legal, accounting, and financial advisory, have historically operated using a pyramidal workforce structure in which large cohorts of graduates and early-career professionals perform research, analysis, drafting, and other developmental work that supports senior-level executives while simultaneously serving as the training ground for future experts. In recent years however, AI is replacing many of these entry-level cognitive tasks. In a controlled study of more than 750 consultants, professionals using AI completed tasks approximately 25% faster while producing higher-quality outputs, reducing the amount of junior labor required to generate equivalent work (Dell'Acqua et al., 2023). Labor market data reflects similar dynamics, with entry-level hiring across financial services declining by approximately 25% in 2025 even as demand for experienced professionals remains strong (Randstad, 2025). The consequence is not simply workforce reduction, but workforce compression. If organizations can maintain or increase output, while reducing the number of individuals entering the profession, the base of the workforce pyramid contracts even though demand for experienced talent remains unchanged. In the near term, this improves efficiency and profitability. Over time however, it weakens the talent-development pipeline through which expertise, institutional knowledge, and professional judgment are cultivated. As fewer early-career professionals gain access to the developmental work that historically produced future managers, partners, and industry leaders, organizations may discover a shortage of experienced professionals capable of performing future work.

Finance and Corporate Functions

Finance, accounting, procurement, and other corporate support departments are also highly exposed functions to AI adoption and workforce compression. These functions are characterized by structured workflows, standardized documentation, and repeatable decision processes, making them particularly susceptible to automation, augmentation, and agentic AI. Tasks such as invoice processing, financial reporting, compliance monitoring, and audit preparation can increasingly be performed faster, cheaper, and at greater scale through AI-enabled systems. Multiple Fortune 500 organizations have linked hiring freezes and headcount reductions within these functions to AI adoption initiatives (Singla et al., 2025). Similar to professional services, the impact is currently concentrated within the early-career roles, while demand for experienced professionals responsible for oversight, interpretation, risk management, and strategic decision-making remains relatively stable. In the current AI market, organizations are able to maintain operational output while reducing the number of employees entering these functions. With time however, this approach narrows the developmental pipeline through which future controllers, finance leaders, procurement executives, compliance officers, and risk professionals acquire the experience necessary to advance. As with other knowledge-intensive professions, the long-term challenge is not simply workforce reduction, but the erosion of the capability-development systems that sustain expertise over time and the workforce pyramid further compresses.

Technology and Software Development

Software engineering provides another clear example of workforce compression rather than workforce contraction. Demand for software continues to grow, yet opportunities for early-career developers have

declined significantly. Between 2023 and 2025, postings for entry-level software engineering positions fell sharply, with highly AI-exposed roles declining by more than 40%, indicating a disproportionate reduction in opportunities for new graduates entering the profession (Revelio Labs, 2025). The primary driver is not a decline in demand for software, but a change in how software is produced. By 2025, AI coding tools generated approximately 27% of all production code across 4.2 million tracked developers, while 93% of developers reported using AI-assisted coding tools at least monthly (DX Research, 2026). Tasks that once required teams of junior developers can increasingly be completed by smaller groups of experienced engineers augmented by AI. Like other functions, organizations can now expand development output, while reducing the number of entry-level professionals required to support the work. While this increases productivity and accelerates delivery timelines in the near-term, over time, it narrows the developmental pathway through which future senior engineers, architects, engineering managers, and technical leaders acquire experience. As the base of the workforce pyramid contracts, the industry risks creating a growing imbalance between the demand for experienced talent and the supply of professionals able to develop into those roles

Customer Service and Operations: AI as the New Front Line

Customer service and contact center operations represents the most visible large-scale deployment of AI as AI agents can now handle a substantial share of routine interactions. Klarna reported that its AI assistant now performs work equivalent to more than 850 customer service agents and has generated substantial cost savings (Klarna, 2025). At the same time, recent evidence is starting to surface, suggesting, AI deployments in this industry are more nuanced, as many firms who originally laid off staff, are now beginning to rehire again as the limits of fully automated service become clearer (Gartner, 2026). Within this function, the workforce effect is now becoming bifurcated: AI systems manage high-volume, standardized interactions, while human workers focus on escalations and complex cases requiring judgment or emotional intelligence. For many organizations, however, the first tasks to disappear are precisely those that once provided entry-level experience.

Manufacturing and Logistics: Automation of Physical Production:

Manufacturing and logistics have experienced automation pressure for decades, but recent evidence shows that AI-enabled robotics and predictive systems are accelerating these trends. Industrial robot adoption continues to expand globally, and workforce studies find “dramatic shifts” in skill requirements as the demand for routine production roles declines, and the demand for technical, maintenance, and systems integration capabilities increases (IFR, 2025; ARM Institute, 2025). The key differentiator here is, unlike other functions, this industry continues to retain structured apprenticeship and vocational training programs, helping to sustain skill development despite automation pressures (OECD, 2025).

Healthcare and Education: Augmentation Over Replacement

At the moment, healthcare and education remain resistant to full automation because of the regulatory constraints, ethical considerations, and the importance of human interactions, judgment and relational trust required. As a result, AI adoption in this industry is primarily augmenting, rather than replacing professionals. AI systems assist with diagnostics, administrative tasks, and personalized learning, while human professionals retain responsibility for interpretation, decision-making, and patient relationships. The industry as a whole has been in a supply and demand deficit, as a result, AI is more of a welcomed solution to offset this deficit. So, while these sectors may still experience shifts in job design and skill requirements, the dominant dynamic right now leans towards role transformation rather than role elimination.

The Common Trend

Despite differences across geographies, demographics, and industries, several consistent patterns emerge from the evidence.

First, AI automates tasks before it automates jobs. However, as increasing numbers of tasks within a role become automated, the cumulative effect often results in workforce restructuring, role consolidation, or eventual role elimination.

Second, workforce impacts are concentrated disproportionately among entry-level and early-career workers. Across nearly every industry examined, AI targets the repetitive, procedural, and developmental work traditionally performed by junior employees. The result is a compression of the workforce pyramid, as organizations reduce the number of individuals entering a profession while maintaining demand for experienced workers.

Third, industries characterized by high levels of human interaction, trust-based relationships, complex judgment, or significant regulatory oversight are currently experiencing AI primarily as an augmentation technology rather than a direct substitute for labor. In these environments, AI tends to enhance human performance rather than replace it outright.

Taken together, these patterns suggest that the most significant consequence of AI adoption is not immediate mass unemployment, but a structural transformation of workforce architecture. By compressing the entry-level layers through which expertise has traditionally been developed, AI risks weakening the talent formation systems that produce future specialists, managers, and leaders. While the pace and severity of this transformation will vary across industries, the underlying trajectory remains consistent: AI is reshaping the pathways through which human capability is formed, creating the conditions for both workforce compression and a broader talent formation fracture.

1.3 Intellectual Atrophy and the Competence Crisis

1.3.1 How Humans Build Expertise And Why AI Is Disrupting this Pathway

The workforce consequences documented in previous sections such as entry-level displacement, workforce pyramid compression, and the talent formation fracture are all visible in market data. Hiring rates can be measured. Job posting volumes can be tracked. The early-career professionals can be identified in payroll records and job posting volumes. However, there is a second category of consequences emerging that has not been discussed with significance or alarm; the erosion of cognitive capabilities from underuse of the mental processes through which competence, judgment, and critical thinking are built because AI is taking over those tasks.

As we know, expertise, in virtually every knowledge-work domain, is not transmitted through instruction alone. It is built through experiential learning and repetition, failure, correction, and the gradual internalization of pattern recognition that comes from doing consequential work under conditions of genuine uncertainty; what we sometimes refer to as the ‘feedback loop’ in the principle of systems thinking. The junior analyst who manually reconciles financial statements is not performing a task the organization values for its own sake. The organization values the analyst who, after two years of

reconciliation, has developed an intuition for where errors concentrate, what anomalies signal, and when the numbers tell a story that the model did not anticipate. The intuition developed is not trainable at the onset, but accumulated over time due to repetition that inherently builds instincts, which then builds the capacity to do harder things that require critical thinking. This is best described in Bloom's taxonomy, a framework established in 1956 and revised in 2001, which categorizes cognition into progressively more complex levels of understanding — from remembering and comprehension, to application and synthesis. Bloom's taxonomy states that higher-order skills are built upon repeated engagement of lower-order processes. Higher-order cognitive tasks like analysis and synthesis do not develop independently, instead they are built through repeated engagement with foundational processes such as knowledge acquisition, comprehension, and application. Therefore, if AI replaces the foundational tasks that serve as the basis for higher-order thinking, the second category of emerging consequences is we lose the ability to develop and strengthen our higher-order cognitive tasks.

This dynamic becomes more significant as AI systems evolve beyond task automation into active participation in higher-order intellectual work. This is demonstrated in emerging "AI Co-Scientist" systems developed by organizations such as Google DeepMind and Future House, which have the capacity to generate hypotheses, synthesize literature, propose experimental pathways, and iteratively refine scientific reasoning in partnership with human researchers. If AI systems increasingly assume responsibility of routine execution, analytical synthesis, conceptual exploration, and probabilistic reasoning, organizations may inadvertently compress both the lower and middle layers of professional development simultaneously. The result is a workforce that becomes progressively more dependent on AI-generated decisions, while having reduced exposure to both the tasks and intellectual struggle that develops judgment, intuition, and expert discernment.

Additionally, the amplification effect of AI with experts versus early career adults is different. Some scholars argue that human cognition has historically adapted to external tools, from calculators to computers, without lasting cognitive degradation. Others contend that AI represents a qualitative shift, because you're outsourcing not just memory and calculation but also reasoning and synthesis, the functions that have traditionally defined high-value human contribution (Jose et al., 2025). In the new future of work scenario, it's possible AI can function as a powerful amplifier and enhance the capabilities of experienced workers, but for less experienced workers, it risks substituting the very processes through which judgment and instincts are developed. The result is a subtle but consequential transformation in how cognition is deployed in professional settings. The bottleneck in knowledge work is no longer the ability to access or process information but how expertise development and critical thinking are cultivated in the work environment.

1.3.2 Cognitive Offloading: The Neurological and Psychological Evidence

Research across neuroscience, cognitive psychology, and education increasingly demonstrates that habitual reliance on AI systems changes how individuals think, learn, encode memory, and sustain the executive functions required for deep reasoning and critical judgment. At the center of this shift is *cognitive offloading* or the experience of *Intellectual Atrophy*™, the tendency to externalize mental tasks to tools rather than perform them internally, thereby weakening human cognition, skills, and judgement in the process.

This phenomenon predates AI however. In 2011, Sparrow, Liu, and Wegner described what is called the "Google Effect," showing that when information is easily retrievable externally, individuals become less likely to encode it into memory. People remember how to access information rather than remembering the information itself. AI extends this dynamic further. Unlike earlier technologies that primarily externalized information storage, AI systems

increasingly offload synthesis, summarization, reasoning, and content generation — cognitive processes that historically required sustained internal engagement.

Recent neurological evidence suggests this shift is already observable at the neural level. In a 2025 MIT Media Lab study, Kosmyna et al used EEG measurements to compare cognitive engagement among participants writing essays with LLM assistance, search engines, or no external tools. Neural connectivity systematically declined as external support increased. Participants writing without AI exhibited substantially stronger and more distributed neural connectivity than those using LLMs, particularly in alpha and theta frequency bands associated with memory retrieval, internally generated thought, working memory, and cognitive control. These findings suggest that participants using AI were not simply retrieving less information from memory — they were actually engaging less in the underlying cognitive processes required to generate, organize, and sustain complex thought. Behavioral outcomes reinforced this pattern. Participants relying on LLM assistance demonstrated significantly weaker memory encoding and lower ownership over their written work. Immediately after completing essays, 83% of LLM-assisted participants could not accurately quote a single sentence they had written. Across the study's four-month duration, AI-reliant participants consistently exhibited lower neural engagement, weaker retention, and less independent compositional integration than participants who completed tasks without AI assistance. In effect, the study provides a neural analogue to the original "Google Effect"; AI appears to reduce engagement not only in remembering information, but in the internal cognitive work required to strengthen neural connections which help us construct ideas themselves.

Parallel findings appear in educational research examining *cognitive load*. Cognitive load theory distinguishes between intrinsic load (task complexity), extraneous load (how information is presented), and germane load — the mental effort devoted to actively processing, integrating, and learning information. In a randomized study of 91 university students, Stadler et al. (2024) they found that students using LLMs during research tasks reported substantially lower cognitive load across all dimensions compared to students using traditional search engines. The largest reduction occurred in germane cognitive load, the component most closely associated with deep learning and schema formation. At first glance, lower cognitive effort appears beneficial, however, the results revealed a more complicated tradeoff. Despite experiencing tasks as easier, students using LLMs produced lower-quality justifications, weaker argumentative rigor, and less evidence-based reasoning than peers using traditional search methods. Importantly, both groups demonstrated similar breadth of perspective, suggesting that AI did not reduce exposure to ideas, but rather reduced the depth with which those ideas were cognitively processed. The difference was not informational access — it was intellectual engagement. In another study performed by Microsoft, where they measured the impacts of Generative AI on critical thinking, 55% and 79% of respondents stated that AI reduced the effort needed to perform core cognitive tasks, and it resulted in the largest reduction in observed comprehension (Microsoft Research, 2025).

Figure 5: Distribution of perceived effort (%) in cognitive activities

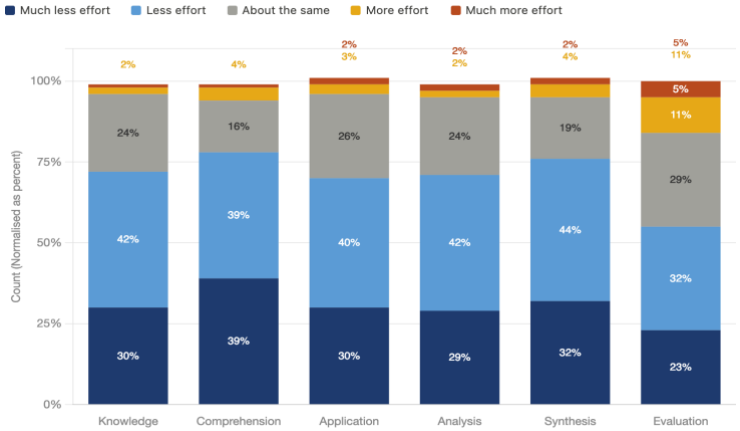

Note: This diagram demonstrates the impact of Generative AI on effort for different critical thinking skills (Microsoft Research, 2025).

Importantly, we are learning that sequencing in the learning process is a large contributor to the depth in which cognitive capabilities are formed or eroded as a result of AI. Individuals who first develop competence independently, then later augment with AI, retain stronger cognitive performance than those who rely on AI from the outset. Any skills developed prior to AI augmentation remain resilient, however skills that were never developed at the onset erode or never form if AI performs the work. This concern is already visible in the aviation industry and they are working towards addressing an alarming concern amongst young pilots. Pilots who over rely on automated systems are experiencing skill degradation Aviation research (2025). As former FAA-recognized chief pilot John Dorcey observed, "When the glass goes dark, you better have skills to fall back on."

The evidence however, is not uniformly negative. Research also demonstrates that AI can be of help and increase creative performance under certain conditions. Doshi and Hauser (2024) asked 300 participants to write short stories either independently or with AI-generated prompts. Independent evaluators consistently rated AI-assisted stories as more novel and useful, with the largest gains occurring among participants who initially demonstrated lower creative ability. In cases where individuals attempt to perform a task outside of their subject matter expertise, where they are less experienced or practiced, AI can raise the performance floor by helping those individuals produce higher-quality outputs. The result however is, while creative quality increased, diversity of expression narrowed. AI in this case, simultaneously elevated individual performance while reducing originality, a phenomenon known as "AI slop."

For organizations, the implications extend beyond productivity metrics. A workforce heavily dependent on AI for analytical and creative work may produce fast outputs, while also gradually losing their originality, diversity, independent reasoning, and developmental depth required for breakthrough innovation and novelty, key ingredients needed for corporate competitive advantage.

Taken together, this emerging body of research suggests that AI's impact on work cannot be evaluated solely through employment statistics or efficiency gains. Its influence on cognition itself, and how individuals learn, reason, remember, and develop novel ideas may prove equally consequential. This add a new paradigm to AI's impact on the workforce since AI is changing how work is performed, but it's also changing how human cognition is developed and sustained over time. Taken together, the neurological and educational findings point toward a risk and pattern: as AI reduces friction in cognitive work, it also reduces the mental effort through which durable understanding, judgment, and reasoning capabilities are formed, which is part of the developmental mechanism through which expertise is built. Over time, the short-term gains in efficiency accompanied by long-term erosion in memory retention, reasoning depth, and independent critical thinking results in Intellectual Atrophy™, the gradual weakening of cognitive capability through persistent dependence on AI and other externalized machine reasoning.

1.4 Conclusion: From Displacement to Developmental Risk

The evidence presented across Part A converges on a single structural conclusion—the impact of AI on the workforce cannot be understood solely through the lens of job displacement. While early indicators point to meaningful disruption at the entry level, with hiring declines, career-ladder compression, and the emergence of a net job deficit window, the more consequential risk lies within the disruption of the pathways by human expertise is formed.

Across industries and geographies, the tasks most susceptible to automation are those that have historically served as the foundation of professional development: drafting, analysis, synthesis, and structured problem-solving. As these tasks are absorbed by AI systems, organizations gain efficiency short term, but simultaneously erode the experiential pathways through which workers build judgment, intuition, and domain expertise. The result is a growing disconnect between the capabilities required at the top of the workforce and the opportunities available to develop them at the bottom.

Additionally, as individuals increasingly rely on AI to perform core components of reasoning and information processing, the effort required to engage in critical thinking declines. While this shift improves the velocity of outputs, it alters the developmental process through which higher-order cognitive capacities are built. What appears as efficiency at the task level may represent degradation at the capability level. Together, these dynamics point to a broader risk in that, the disruption of AI is not simply a shortage of jobs, but a weakening of human capability that sustains long-term economic and organizational performance. The implications extend beyond labor market outcomes.

Organizations that optimize for output while eliminating developmental pathways may find, over time, that their pipeline of capable talent has narrowed. Economies that experience widespread entry-level displacement may face not only unemployment, but the structural exclusion of entire cohorts from the processes through which expertise is built. In this sense, the challenge posed by AI is not merely economic, it is developmental, institutional, and intergenerational. These dynamics raise a critical question, if AI is restructuring not only work but the pathways through which human capability is developed, what must be deliberately redesigned to preserve and sustain that capability in an AI-augmented economy?

PART 2:

RECONSTRUCTING THE WORKFORCE IN THE AGE OF AI

2.1 The Case for Immediate Intervention

2.1.1 Reimagining Human Intelligence in an AI-Augmented World

As we learned, higher-order capabilities do not emerge independently of foundational work. Expertise is built through repetition, practice, error correction, and the gradual accumulation of domain-specific knowledge; and the ability to exercise sound judgment often depends on years of experience performing the very tasks now being delegated to AI. If individuals no longer develop the foundational competencies that underpin expertise, the higher order capabilities organizations seek may never fully mature, posing a risk to develop other cognitive and highly critical functions.

In the future of work with automation, augmentation, and agentic AI, there's consensus that leaders will need to continue to lead companies, navigate market conditions and economic ambiguity; and what creates exceptional leaders is not ability to recall facts, but their capacity to navigate ambiguity, integrate disparate perspectives, make principled decisions, and recognize what others—including increasingly sophisticated AI systems—fail to see. Ironically, these are the very capabilities most at risk if AI absorbs the experiential work through which these capabilities are developed. Global workforce data reinforces this reframing and the World Economic Forum projects that 59% of the global workforce will require reskilling or upskilling by 2030, with 11 in every 100 workers unlikely to receive it, placing them at medium-term risk of redundancy (WEF, 2025). Additionally, these new skills consistently prioritize both technical and functional skills, where AI literacy and technological fluency sit alongside analytical thinking, creative thinking, resilience, and adaptability. Without the pathway to execute, test, and learn these skills we create a further gap in developing talent needed to support our AI future.

McKinsey's brain capital framework, developed with the WEF and published in January 2026 also determined that the new economic value of the future have these two interdependent assets intertwined; technical skills with interpersonal and self-leadership abilities that enable people to adapt, contribute, and sustain high-level performance over time. As AI reshapes work, the report argues, competitive advantage will increasingly hinge on how well organizations combine human and machine strengths, and that cannot happen without deliberate investment in brain evolution, cognition, and training of talent (McKinsey Health Institute & WEF, 2026).

When further exploring the workforce implications, we discovered the impacts extend beyond cognitive and productivity metrics. Jobs provide income, but at its core it also provides identity, mastery, and social recognition. When AI compresses or automates significant portions of intellectual contribution, it does not just threaten physical labor, it threatens cognitive authority, self-purpose, and aspects of our identity, the foundation of professional identity in knowledge-work economies (Brooks, 2024). This produces what might be called 'Relevance Anxiety,' a persistent uncertainty about whether one's skills, experience, and judgment will retain value as AI systems become more capable. 'Relevance anxiety' is particularly acute for two groups, early-career workers denied the entry points through which professional identity is formed, and mid-career professionals whose domain expertise is being partially automated and obsolete before they have developed the breadth to pivot into supervisory or strategic roles.

At the societal level, these dynamics risk producing a bifurcated workforce: a relatively small group of highly leveraged orchestrators and domain specialists whose cognitive authority is amplified by AI, and a much larger group experiencing diminished agency, relevance, and Intellectual Atrophy™. Without deliberate institutional intervention, AI will exacerbate economic inequality and cognitive and psychological stratification, a stratification that is harder to measure, slower to become visible, and more difficult to reverse than the employment statistics that currently dominate policy attention.

Additionally, without deliberate investment in human capital, AI adoption may erode precisely the cognitive capabilities that remain uniquely and irreplaceably human. Early-career workers who outsource reasoning before developing domain expertise, intuition and pattern recognition will not naturally develop the oversight capacity that the pending 'Oversight Economy' requires. Organizations that optimize purely for output will produce employees who are efficient in the short term, and cognitively fragile over a longer horizon.

The recommendations that follow to begin the process of combating these forces and threats on human capital operates on two fronts: talent development and workforce architecture in order to preserve capability formation within firms; and policy interventions, to realign the incentives that govern AI adoption across the economy.

2.2 Blueprint for Deliberate Redesign for Businesses

In the previous section, we established a central workforce risk, outside of job loss, and that is the compression of the workforce pyramid, and the very breakdown of the system that helps build human expertise needed to oversee the AI ecosystems of the future. As entry-level work disappears, and work becomes more automated, a talent formation fracture occurs, posing a threat to how organizations build the expertise needed to govern our future AI ecosystems and the modified cognitive skills needed to stay relevant in the age of AI.

As such, this section begins the process of outlining strategic planning and intervention vehicles to help combat the agitation AI is creating in the workforce. The frameworks that follow synthesize findings from six field studies on AI adoption and workforce outcomes: MIT Project NANDA (Challapally et al., 2025), Humlum & Vestergaard (2025), Brynjolfsson et al. (2018, 2025), Acemoglu (2022), Dell'Acqua et al. (2025), and McKinsey AI Survey (2025). Their findings conclude AI deployment without workforce redesign produces no measurable organizational transformation, and workforce architecture redesign is what separates the 6% of firms capturing real AI value from the 94% that are not. These studies contribute to the blueprint, framework, and recommendations made in the subsequent sections.

2.2.1 The New Talent Taxonomy

We first begin by dissecting and rebuilding the 'Talent Taxonomy' of any given role. Meaning, how should a role evolve within an organization to stay relevant, and provide capital value in the age of automation, augmentation, and agentic AI. The method is a structured synthesis of the most established peer-reviewed talent model in management literature (the 'T-shape') and examining post-2020 evidence on how AI is challenging that model.

The M-Shaped Talent Taxonomy Model

For more than twenty years, workforce development, corporate hiring, and organizational design have all relied on a single model, the T-shape talent taxonomy model. The vertical bar in the 'T' represents deep expertise in one specialty, while the horizontal bar represents broader skills like communication, collaboration, and awareness of how things connect, that allows specialists to work across teams. The model became standard after Hansen and von Oetinger (2001) wrote about it in the Harvard Business Review, and for nearly two decades it has been cited across a variety of academic works, including Demirkan and Spohrer (2015) who extended the work into representing skills of the modern workforce; individuals with "deep expertise in at least one field" who can couple that skill with systems thinking, allowing them to thrive in environments that require communication and work across disciplines, sectors,

and cultures. The 'T-shape' model however, is only applicable in stable and constant environments with clear roles and boundaries, environments where career ladders are relatively predictable, workforce pyramid is a dominant structure, and technology is leveraged as a tool, not a coworker. However, in the age of AI, where automation, augmentation and agentic AI are disrupting the workforce and the core tasks, subtasks, and work itself, the quarter-century long T-shaped skills model is being challenged.

First, the value of deep expertise on routine work is shrinking. Brynjolfsson et al. (2025) studied 5,172 customer-support agents at a Fortune 500 company in a Quarterly Journal of Economics paper. They found that AI assistance produces about a 15% productivity gain on average, but the gains are uneven: "Less experienced and lower-skilled workers improve both the speed and quality of their output while the most experienced and highest-skilled workers see small gains in speed and small declines in quality." The researchers interpret this as AI passing tacit best practices from experts to less-skilled workers, closing the gap between novices and experts on routine expert work. Therefore, the value of having a single deep specialty or expertise in one field for routine work is shrinking.

Second, the work left for human judgement is concentrated on tasks AI cannot evaluate without a human expert in the loop. Lebovitz et al. (2022) ran an ethnographic study of diagnostic radiologists at a major US hospital, published in Organization Science. They describe how doctors handled the "hindering effects of opacity" they encountered when using AI tools. Complementary research in MIS Quarterly shows that even the labels used to train medical AI come from expert tacit judgment (Lebovitz et al., 2021). The research outlined in both studies indicate that domain expertise does not lose its value within AI ecosystems, however the way in which that expertise is applied is changing. The application of expertise and knowledge changes from producing the answer to judging the AI-generated answers. This new evolution is a different skill set and requires different supporting skills to be useful at scale.

Third, the workflow becomes the unit of analysis, not the tasks within the workflow. Dell'Acqua et al. (2025) introduce the idea of a "jagged technological frontier… [where] AI assistance improves performance for some tasks but worsens it for others" and therefore the entire workflow needs to be assessed when wanting to extract business value from AI. Combined with the McKinsey finding that workflow redesign is the single largest predictor of AI-attributable profit out of 25 attributes tested (McKinsey, 2025c), workflow design becomes a worker-level skill.

Taking the results from above, our thesis illustrates the future of work is requiring roles to evolve into four (4) interconnected depths resulting in a new talent taxonomy model, the 'M-shape' model, replacing the historic 'T-shape' model. Each of the four (4) pillar matches a skillset observed as necessities based on the post-2020 research. The four pillars are: domain expertise, AI literacy, systems thinking and workflow design, and human judgment with contextual reasoning.

Pillar 1: Domain expertise is the skill anchor, possessing real depth in a specific field. The kind of mastery the T-shape vertical bar describes. However, in the age of AI, the job of expertise shifts from producing the answer to judging whether the answer of an AI application is accurate and relevant based on the situation and context. Without genuine expertise, the "human expert in the loop" cannot reliably distinguish fact from hallucination, correlation from causation, or sound reasoning from confident-sounding error. This dynamic emerged in a 2025 customer-support field experiment (Erik Brynjolfsson and colleagues). They discovered AI created the greatest performance gains for novices and early-career workers by elevating their output closer to expert levels, while experienced experts saw declines in quality when over-relying on AI. AI's greatest value may not be replacing expertise but compressing the distance between novice and expert performance — while simultaneously making true human judgment, oversight, and contextual discernment more strategically important than ever.

Pillar 2: AI literacy is emerging as a foundational competency for the future of work. The workforce needs to have a practical understanding of how AI systems behave, where they fail, and how to guide them toward more reliable outcomes. This includes recognizing the probabilistic nature of AI outputs and knowing how prompts, constraints, context windows, escalation protocols, and validation checks shape the behavior of AI agents. AI literacy at this level should not be confused with software engineering expertise or advanced technical certification. Rather, it resembles financial literacy for managers. A manager does not need to be an accountant to lead a business effectively, but one does need enough fluency to interpret, question, and challenge financial statements with legitimate critical authority. Duri Long and Brian Magerko, define AI literacy as "a set of competencies that enables individuals to critically evaluate AI technologies, communicate and collaborate effectively with AI, and use AI as a tool online, at home, and in the workplace" (Long and Magerko, (2020). Their framework positions AI literacy not as a narrow technical skill, but as a broader form of societal and organizational fluency. Daron Acemoglu and colleagues show that firms adopting AI technologies are simultaneously changing the skill requirements of many traditionally non-AI roles, suggesting that AI fluency is diffusing beyond technical occupations into general managerial and operational work (Acemoglu et al., 2022). Likewise, Tatyana Alekseeva and colleagues, writing in Labor Economics, identify a measurable wage premium associated with AI-related skills across a wide range of occupations, not just within engineering or software development (Alekseeva et al., 2021).Together, these findings suggest AI literacy is evolving into a baseline workforce capability: not because every worker must build AI systems, but because an increasing number of workers will be expected to supervise, interrogate, collaborate with AI systems, and make judgments alongside them.

Pillar 3: Systems thinking and workflow design are emerging as the second critical pillar in skill development. The skill of the future is no longer focused on performing an isolated task, but understanding how AI agents, human teams, data pipelines, governance mechanisms, and organizational processes interact as an integrated system, a unified ecosystem, and then designing work that allow those components to operate effectively together. This perspective aligns with the task-based framework proposed by Daron Acemoglu and Pascual Restrepo (Acemoglu and Restrepo (2019). The research concludes the economic outcome of AI adoption depends less on the capability of the technology and more on how work itself is redesigned around it. Workers consistently struggle to determine where AI improves performance versus where it degrades it at the workflow level, suggesting that orchestration itself is becoming a distinct professional competency (Dell'Acqua et al. (2025)). Similarly, Diane Bailey and Stephen Barley Bailey and Barley (2020), argue that technological outcomes are shaped primarily by the surrounding work-system context rather than by the technology alone.

Pillar 4: Human judgment, ethical reasoning, and contextual interpretation are becoming more valuable as AI systems become more capable. These capabilities involve making decisions when there are no clear rules, recognizing the ethical and social implications embedded within seemingly technical outputs, and applying empathy, trust, and contextual awareness in situations where human relationships shape outcomes. Many of the highest-stakes organizational decisions across industries such as healthcare, legal, supply chain, manufacturing contain ambiguity that cannot be fully reduced to structured logic or historical pattern recognition. The foundational argument for this distinction comes from David Autor, whose work on "expert thinking" explains that humans retain an advantage in tasks involving abstract reasoning, unstructured problem-solving, tacit knowledge, and complex interpersonal interaction because such activities cannot easily be codified into explicit rules for machines to execute (Autor, 2015). More recent organizational research extends this perspective into the AI era. Chris Anthony and colleagues argue human contribution increasingly shifts toward judgment under ambiguity, contextual interpretation, and the resolution of conflicts between technical optimization and human consequence (Anthony et al., 2023). Knowing when to trust AI systems and when to override them, is itself a learnable capability, and that poor calibration between over-trust and under-trust is one of the primary reasons human-AI teams fail (Glikson and Woolley, 2020). Michael Polanyi theory of tacit knowledge suggests that a significant

portion of human expertise is embodied, contextual, experiential, and socially situated rather than fully articulable into a workflow. To the extent that AI systems are trained primarily on what can be explicitly documented and encoded, there remains a category of human judgment and lived understanding that is difficult to replicate computationally (Polanyi, 2009).

Figure 7: The T-shaped and M-shaped professional talent models

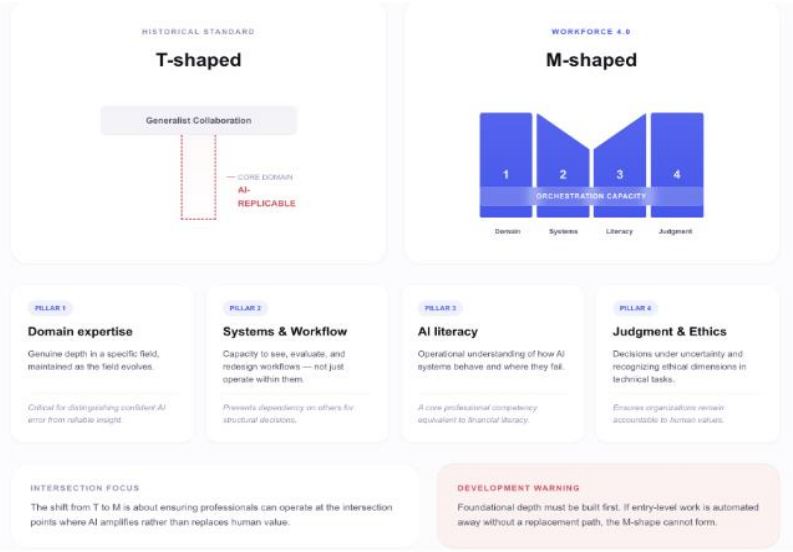

Note: This figure illustrates the T-shaped and M-shaped talent taxonomies side by side. The T-shaped model (left) represents the dominant pre-AI era skill model: deep expertise in a single domain bridged by broad collaboration skills. The M-shaped model (right) represents the four-pillar architecture proposed in this research, requiring the workforce that constitutes the future of work in the AI era to have 4 dominant pillars of skills: domain expertise, systems thinking and workflow design, AI literacy, and human judgment.

Organizational Implications When Shifting from T-Shaped to M-Shaped Model

It's important to note, the transition from T-shape to M-shape capability structures does not imply the future of work means all workers must become technologists or AI engineers. Rather, it reflects a broader shift in the cognitive and organizational demands of work under conditions of AI augmentation. Increasingly, workers are expected to execute tasks, but also to interpret AI-generated outputs, evaluate their reliability, understand the contextual conditions under which they were produced, and exercise judgment in situations where the system's recommendations are incomplete, ambiguous, or potentially incorrect. As AI systems become embedded within operational workflows, the value of labor shifts from routine execution, and more toward supervision, contextual reasoning, orchestration, and governance.

There are two organizational implications follow as a result of this transition. First, the M-shape skill requirements cannot be developed through conventional training programs alone. Classroom instruction and static upskilling modules are insufficient because the competencies in question are fundamentally experiential and situated within live operational systems. Workers develop these capabilities by participating directly in AI-mediated workflows, making decisions alongside AI systems, and learning how human and machine judgment interact under conditions of uncertainty. AI-exposed firms do not merely add AI training to existing organizational structures; rather, they reorganize the structure of work itself, altering task composition, decision rights, and workflow design (Acemoglu et al., 2022). The implication is that capability formation in the AI era depends less on isolated technical instruction and more on redesigning the environment in which work occurs and learning on the job.

Second, domain expertise remains the foundational layer upon which the remaining M-shaped competencies depend. If organizations continue reducing entry-level hiring or automating away junior-level cognitive work, the developmental substrate required to produce future expertise erodes over time. This reflects, at the organizational level, the broader "talent formation fracture" described in previous sections at the labor-market scale. Across most knowledge-intensive professions, expertise is historically accumulated through repeated exposure to uncertainty, iterative problem solving, supervised correction, and gradual pattern recognition developed through direct engagement with operational work. Tasks such as drafting, reconciliation, preliminary analysis, and first-pass research have traditionally functioned as developmental mechanisms through which junior workers internalize judgment and tacit knowledge. When these activities are delegated primarily to AI systems, organizations will remove not only labor inputs, but also the experiential pathway through which expertise itself is formed. Research suggests that while AI systems can elevate novice performance toward expert-level outputs, the developmental trajectory through which expertise is acquired becomes structurally different from pre-AI learning pathways (Brynjolfsson et al., 2025). Similarly, Behnam Beane, in an ethnographic study of robotic-surgery training within leading US teaching hospitals, observed that trainees lost access to critical forms of skill acquisition once intelligent systems assumed core procedural tasks (Beane's, 2019).

Taken together, these findings suggest that the central challenge for organizations and employees now, in the AI-era transformation is not simply technological adoption, but more so, workforce redesign that needs to evolve and complement the technology being adopted simultaneously. Organizations must go through a series of steps in their workforce planning strategies that encompasses redesign their workforce architecture, redefine how work should be structured, how roles and tasks evolve, how expertise will be developed and how judgment authority is distributed between humans and intelligent systems. The transition from T-shape to M-shape talent models therefore represents more than a skills evolution; it reflects a broader shift in the architecture of work itself; and the organizations that succeed will not be those that automate the fastest, but those that most effectively redesign the relationship between human capability formation, workflow structure, and AI-enabled decision systems.

2.2.2 Role Archetype Considerations for the AI Era

Following the typology-as-theory-building approach described by Edward Doty and William Glick, typologies function as explanatory theoretical models that define ideal-type configurations, they specify relationships among categories, and generate testable predictions (Doty and Glick, 1994). Therefore, as part of the strategy for workforce planning in the age of AI, we recommend a redesign in role architecture, which compliments the M-shape talent taxonomy model and provides a framework to support workforce preparation for Workforce 4.0. This framework is called the OOA framework. The OOA framework lays out three role configurations and examines the interdependencies among them, along with their implications for productivity, governance, and skill formation. The analysis proceeds in three stages: first, it synthesizes research and literature into why a typology; second, it explores the typology's structure and theoretical grounding; and third, we apply the framework to a representative role to show case its application in real life using O*NET task descriptions.

Why a Typology

As we know it, the task composition of knowledge workers is undergoing structural transformation under AI adoption. Across four independent approaches to measuring occupational AI exposure (Brynjolfsson et al., 2018; Felten et al., 2021; Webb, 2020; Eloundou et al., 2024), consistent finding emerge: most occupations are composed of a heterogeneous mix of AI-suitable and AI-unsuitable tasks. The implication is that AI adoption does not simply automate entire jobs but instead reshapes the internal allocation of tasks and activities within them. Moreover, the research consistently states, ROI on AI is only extracted if

the technology deployment is complimented with a redesign of tasks and subtasks that make up a role, to account for the new architecture of work in the age of automation, augmentation, and agentic AI. Within this context, a role typology serves as an organizing mechanism for the structural changes required to support evolving patterns of human–AI collaboration and the broader reconfiguration of knowledge work in the age of AI.

The Three Archetypes

Within the topology structure, we've identified three (3) archetypes — the Operator (O), the Orchestrator (O), and the Artisan (A). We will refer to this as the OOA framework. Each of the three archetypes has a consistent analytical structure: a working definition, representative forms of work, and a strong empirical anchor. However, two clarifications are important.

First, the archetypes describe modes of work within AI-augmented organizations rather than fixed occupational titles. Individuals may move dynamically across archetypes depending on the task, workflow, and decision context involved. For example, a physician may operate within the Artisan archetype while making a complex diagnostic judgment but shift into the Operator archetype when supervising an AI system conducting routine image analysis. Similarly, a product manager may alternate between Orchestrator and Operator functions depending on the work they are performing for the company. The typology therefore partitions work activities and functions, not workers themselves.

Second, the OOA framework should not be interpreted as a hierarchical ranking of labor value. The three archetypes represent interdependent organizational functions that coexist within an integrated architecture. Orchestrators design and govern workflows, Operators manage and supervise AI-mediated execution, and Artisans lend deep expertise, contextual reasoning, and tacit judgment. The effectiveness of the framework depends on maintaining the coordination among the three modes. In this sense, the typology and the organizational architecture are inseparable components; the typology defines human contribution within an AI ecosystem, while the architecture specifies how those contributions are structurally distributed, coordinated, and governed across workflows. The OOA framework therefore functions as a classification scheme for work, but also as an organizational design model for allocating decision rights, capability development, and human–AI interaction within AI-enabled enterprises.

Figure 8: The OOA Framework

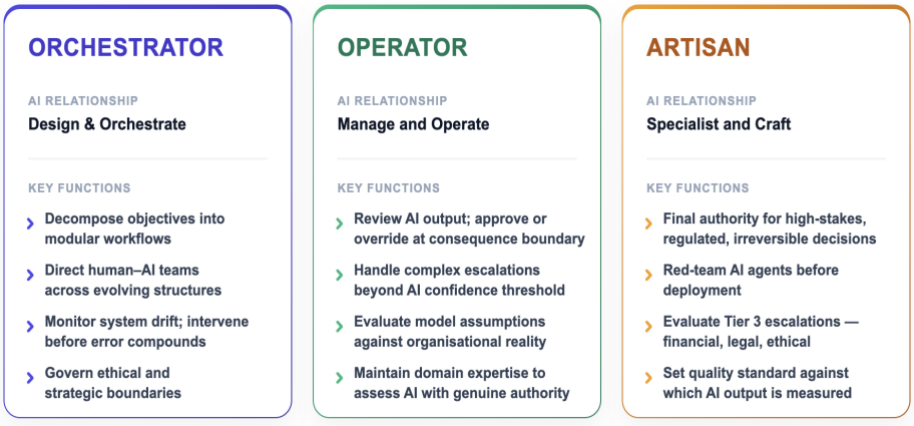

Note: This figure introduces the Orchestrator-Operator-Artisan (OOA) Framework, a proposed role architecture for organizing human contribution and work in AI-augmented environments.

The *Orchestrator* serves as the systems architect of the AI-enabled organization. Unlike traditional managers, whose primary responsibility is coordinating people within established organizational structures, the Orchestrator is responsible for designing and continuously optimizing the division of labor between humans and intelligent systems. The role centers on decomposing business objectives into workflows, determining which activities should be automated, augmented, or retained as human-led work, and continuously refining those decisions as technologies, capabilities, and organizational priorities evolve. Orchestrators oversee the interconnected ecosystem of AI agents, data pipelines, business processes, governance mechanisms, and human expertise required to produce reliable outcomes at scale. They are responsible for ensuring that AI systems remain aligned with business objectives, that data flows support consistent performance, and that evolving technologies are integrated into operational environments without introducing unnecessary risk, inefficiency, or capability degradation. This role emerges in response to one of the central themes in this thesis in that, extracting value from AI is not primarily a technology problem, but a workflow-design problem. For example, BCG examined the use of AI among 758 consultants and found that one of the greatest determinants of performance was not access to AI itself, but the ability to recognize where AI creates value, where it introduces limitations, and where human judgment remains essential (Dell'Acqua et al., 2025). Determining the appropriate boundary between human and machine contribution therefore becomes a distinct organizational capability, and the Orchestrator is responsible for managing this boundary. In doing so, the role operationalizes what Raisch and Krakowski (2021) describe as the automation-augmentation paradox: the continuous balancing of activities best performed by machines and those best retained by humans. As organizations become increasingly dependent on AI-enabled workflows, the ability to design, govern, and adapt this division of labor becomes a critical source of organizational performance and competitive advantage.

The Operator serves as the execution and oversight layer of the AI-enabled organization. As AI systems assume responsibility for increasing volumes of routine, repeatable, and rules-based work, human involvement shifts from performing tasks to supervising outcomes. The Operator's primary responsibility is not to execute the work itself, but to monitor AI-generated outputs, manage exceptions, identify failures, and intervene when situations require context, judgment, or accountability beyond the capabilities of the system. This archetype emerges from a fundamental characteristic of AI deployment: while intelligent systems can automate large portions of operational work, they remain constrained by the quality of their training data, the assumptions embedded within their models, and their inability to fully interpret organizational context. As a result, organizations continue to require individuals capable of evaluating outputs, identifying anomalies, resolving exceptions, and determining when human intervention is necessary. The Operator functions as the bridge between automated execution and human judgment, ensuring that AI-generated decisions remain aligned with business realities and operational objectives. Evidence of needing this archetype is increasingly visible across industries. For example, research examining customer-support environments found that AI assistance improved productivity by approximately 15 percent while simultaneously reducing the volume of interactions requiring direct human involvement (Brynjolfsson et al., 2025). Similar patterns have been observed across operational functions as algorithmic systems assume responsibility for routine activities and human work shifts toward verification, supervision, and exception management (Pachidi et al., 2021). Rather than eliminating human participation entirely, AI changes the nature of human contribution. Within the OOA framework, the Operator represents the human supervisory layer responsible for maintaining accountability within increasingly automated workflows. This aligns closely with what Shrestha et al.'s (2019) describe as delegated decision architectures, in which routine decisions are performed by AI systems while humans retain responsibility for oversight, escalation, approval, and final accountability. As AI adoption expands, the ability to effectively supervise intelligent systems may become as important as the ability to perform the underlying work itself.

The Artisan represents the domain expert within the AI-enabled organization. Unlike the Operator, whose role is responsible for supervising AI-mediated execution and the Orchestrator designs the human–AI

operating model, the Artisan performs the work that remains resistant to standardization, prediction, and automation and engages in forms of reasoning where standardized pattern recognition alone is insufficient. In radiology, for example, the Operator may oversee and approve AI-assisted interpretation of routine imaging, whereas the Artisan evaluates rare presentations, ambiguous findings, or diagnostically complex patients whose histories fall outside the statistical regularities on which the model was trained. In legal practice, the Operator may review AI-generated contracts or routine filings, while the Artisan develops novel legal arguments, interprets unsettled areas of doctrine, creates policy change, or litigates cases requiring strategic persuasion and contextual interpretation before a court. This archetypes core value is to operate in environments characterized by ambiguity, uncertainty, contextual complexity, and incomplete information, where success depends on technical competence, judgment, experience, and the ability to navigate situations for which no predefined rule set exists. The distinguishing characteristic of the Artisan is the application of domain-specific reasoning in circumstances where statistical prediction alone is insufficient. Across professions, this includes evaluating novel situations, resolving conflicting information, interpreting context, making decisions under uncertainty, and generating new knowledge, insights, or solutions. As a result, the Artisan's relationship with AI is fundamentally augmentative rather than substitutive. AI may expand the scale, speed, and informational reach of the work, but value creation remains dependent upon the human expert's ability to integrate contextual understanding, tacit knowledge, institutional experience, and professional judgment. Furthermore, this role becomes increasingly important as AI systems assume responsibility for routine prediction, classification, drafting, and pattern recognition. Research suggests that AI systems are typically trained on codified representations of expertise rather than the full breadth of reasoning used by experts in practice (Lebovitz et al., 2021). Similarly, studies examining AI-enabled hiring systems found that organizations ultimately reintroduced human expertise because machine learning systems alone proved insufficient for interpreting complex social, contextual, and situational variables (van den Broek et al., 2021). These findings suggest that as organizations automate standardized execution, they simultaneously increase their dependence on individuals capable of resolving uncertainty when statistical prediction becomes inadequate. Within the OOA framework, the Artisan serves as the primary source of judgment, innovation, and adaptation. The role not only addresses the limitations of AI systems, but also contributes to the continued development of organizational knowledge, new ideas, and emerging practices that AI systems cannot independently generate. In this sense, the Artisan operates as a stabilizing function within the broader human–AI ecosystem, integrating contextual awareness, ethical reasoning, institutional memory, and domain expertise in situations where the consequences of ambiguity, error, or misinterpretation are significant. As AI absorbs increasing amounts of routine cognitive work, the ability to exercise expert judgment under conditions of uncertainty becomes an increasingly valuable and differentiated human capability.

It's important to note, the Orchestrator, the Operator, and the Artisan are not hierarchical in the usual corporate sense. They are functional, value-generating roles that reflect different ways humans contribute when AI is in the picture, and they need different kinds of accountability around them.

The Application of the Archetypes in the Workforce

The three archetypes are not independent roles to staff separately, nor are the interdependencies symmetric. They are interdependent functions an organization must establish as part of their talent planning strategy. The Orchestrator depends on Operators with real operational substance, to design how the workflows operate and what gets routed to humans versus AI. The Orchestrator also depends on the Artisans for their domain expertise and to set quality standards that anchor what good AI outputs looks like in functions where AI's quality bar is set by expert craft (radiology, complex investment work, legal arguments) rather than by codified rules or regulatory checklists. The Operator and the Artisan in turn depend on Orchestrators for the workflow design that puts AI at the right tasks at the right thresholds. Without that design, neither the Operator's checkpoint authority nor the Artisan's domain craft is matched against problems where AI augmentation pays off. Therefore, the dependencies across the archetypes

form a network rather than a closed cycle, with the Orchestrator positioned at the center because it provides the coordination and strategic oversight that make the other roles coherent.

Figure 9: The OOA Architecture of Interdependence

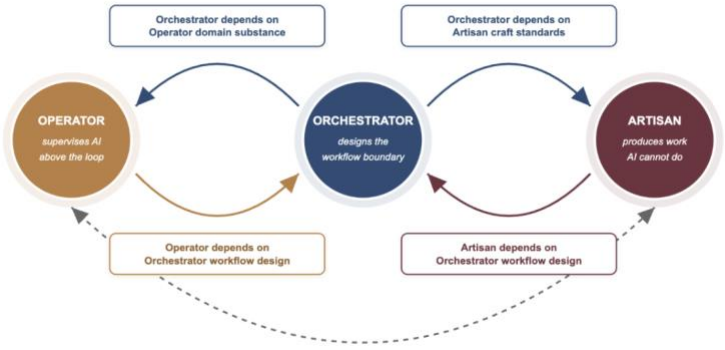

Note: This figure illustrates the interdependence between the Orchestrator, Operator, and Artisan archetypes within the OOA Framework. Rather than functioning as isolated roles or a linear hierarchy, the framework conceptualizes human work in AI-augmented organizations as an interconnected system in which each archetype depends on the others to sustain effective execution, governance, adaptation, and expertise.

The M-shape talent taxonomy model and OOA operate at different units of analysis but with interdependencies. The M-shape model establishes the capabilities needed for knowledge workers to function in an AI-augmented workforce. The OOA framework establishes the roles each knowledge worker plays and how the work gets distributed. The roles knowledge workers play and the capabilities they must possess are not to be confused as one in the same. For example, an Operator with M-shaped skills, possess the capabilities needed to be the qualified human checkpoint in AI ecosystems. With the depth of the M-shaped model, the Operator has the qualifications and authority to carry substantive evaluative power, including the authority to override AI output, escalate decisions, and trigger workflow redesign upward to Orchestrators. Without the depth and span of skills the M-shaped model suggests, the Operator cannot be effective. While each archetype needs all four pillars presented in the M-Shape talent taxonomy model, the concentration and mix of each pillar will vary by archetype.

Archetype	**Primary pillars**	**Secondary pillars**
Orchestrator	Systems thinking and workflow design AI Literacy	Domain expertise Human judgment and ethics
Operator	AI Literacy Human judgment and ethics	Domain expertise Systems thinking
Artisan	Domain expertise Human judgment and ethics	Systems thinking AI literacy

The Orchestrator designs and revises the human-AI workflow boundary, which leans on systems thinking and an understanding of what AI can and cannot do across the workflow. The Operator supervises AI execution in real time, which leans on AI literacy and judgment to weigh consequences against AI confidence. The Artisan does the irreducibly expert craft work that AI cannot perform end-to-end, which leans on deep domain expertise and on human judgment that define thresholds for good output.

Real-World Example: Search Marketing Strategist

To demonstrate how the OOA framework operates in practice, this section applies the model to the occupation of Search Marketing Strategist (O*NET 13-1161.01). The occupation was selected for three reasons. First, it represents a domain in which AI disruption is already observable and empirically documented. Field experiments by Hartmann et al. (2025) and Lee et al. (2025) have already tested AI-augmented marketing workflows in live campaigns at scale, providing real-world evidence of workflow restructuring under AI integration. Second, the role is economically significant rather than a niche or experimental occupation. Search Marketing Strategist is a detailed occupation within the broader category of Market Research Analysts and Marketing Specialists (SOC 13-1161), a category employing approximately 941,700 U.S. workers at a median annual wage of $76,950, with projected growth of 7% or higher through 2034 according to the U.S. Bureau of Labor Statistics (U.S. Bureau of Labor Statistics, 2025, via O*NET). Third, the occupation has already been evaluated within published AI exposure research, particularly the task-level exposure framework, making it suitable for applying OOA's task-to-governance assignment methodology (Eloundou et al., 2024).

Search Marketing Strategist serves as a representative Operator archetype occupation. The Department of Labor's O*NET database lists 36 distinct tasks for Search Marketing Strategists. These tasks are not grouped into families, the 36 tasks are outlined in a flat list, with roughly half of the 36 tasks documented as already performable by AI today. We however, have grouped the tasks into four (4) families in order to create an analytical layer that makes visible how AI exposure is distributed unevenly across different forms of work within a single occupation. Additionally, we ensured that the family assignments supports deductive thematic coding, following the protocol established by Boyatzis (1998) and operationalized by Braun and Clarke (2006) (see Appendix B).

The 36 tasks have been grouped into the following four families:

1. *Content production and optimization:* Tasks which create media artifact (text, image, page, and creative) intended for downstream use and optimization of the artifacts against measurable engagement criteria. Examples include: creating digital assets, optimizing content for SEO, preparing creative prototypes.
2. *Measurement, analytics, and reporting:* Tasks that aggregate, transform, or interpret performance data into outputs (dashboards, reports, models) used by other roles. Examples include: collecting and interpreting web metrics, building dashboards, modeling campaign economics.
3. *Campaign execution and operations:* Tasks that manage marketing plans against budgets, schedules, and performance targets. Examples include: A/B tests, running paid ad placements, email and social campaigns.
4. *Strategy, judgment, and external relationships:* Tasks that decide what the function should and should not do, negotiation or representation outside the function. Examples include: proposing campaigns to executives, evaluating emerging media, negotiating with vendors, managing regulatory and brand risk.

Each of the 36 O*NET tasks was assigned to the task family whose definition most closely reflected the underlying nature of the work. Figure 10 maps these tasks across the four families and highlights which activities can already be performed end-to-end by AI. Beyond creating a more structured classification

system, this approach illustrates a critical limitation in how current organizations often assess the workforce. Organizations frequently evaluate automation at the role level, implicitly treating occupations as either replaceable or non-replaceable. In practice, however, AI does not automate jobs; it automates specific tasks within roles.

Figure 10 shows how tasks within the same occupation vary considerably in their susceptibility to automation. Some are repetitive, rules-based, and prediction-oriented, making them highly automatable; while others require contextual interpretation, organizational judgment, domain expertise, or human interaction. It reveals how work is redistributed as AI is introduced, and identifies where automation, augmentation, and human expertise coexist. More importantly, it shifts the analysis away from simplistic assumptions of role elimination and toward a more accurate understanding of how AI reshapes work at the task level. The result is not the disappearance of entire occupations, but the reallocation of activities within them, creating new opportunities for automation, augmentation, and human specialization.

Figure 10: The 36 tasks of a Search Marketing Strategist, by family and AI exposure.

*Note: This shows a heat map of all 36 O*NET-listed tasks for Search Marketing Strategists (SOC 13-1161.01), grouped into four task families. Tasks in red are E1-rated by Eloundou et al. (2024), performable end-to-end by current LLMs and nearly fully automated tasks. Tasks in amber are E2-rated, indicating some aspects of the task itself can be performable through AI if embedded into specialized marketing software, but the task itself still requires human context and judgement. Tasks in green are E0, currently not exposed to AI today. Of the 36 tasks, 10 are fully automated, 1 cannot be automated, and the remaining 15 are tasks augmented by AI but not fully replaceable as of yet.*

Using the approach of decomposing a role into discrete tasks, identifying which tasks can be fully automated, partially augmented, and which ones remain fully intact, you can view the overall composition of tasks that make up a role and then interpret its archetype so that you can best cultivate, evolve, and curate your workforce in the human-AI division of labor. In the case of the Marketing Strategist, this role is evolving into an Operator, which means when designing your marketing organization in the age of AI, you'll need to identify employees who can function as the Orchestrators to support the Operators work, and the Artisan, who sets the standards by which the Operator executes against. However, if you find yourself have a team of many Operators with very few or no Orchestrators or Artisans, that's where your workforce design begins. You assess talent for their archetypal competency, reallocate individuals to create the dynamic of the OOA Framework, and begin cultivating this organizational design across all functions.

Boundary Conditions

It's important to note, this framework does not apply to, without significant adaptation, in three cases: (a) predominantly physical labor (assembly-line, skilled trades, in-person care), this is better addressed by the robotics-and-automation literature; (b) work where regulation already prescribes a specific human-AI division of labor (certain FDA 510(k) clinical workflows, SR 11-7 credit-risk decisions, certain EU AI Act high-risk applications), where the tiers are already determined by regulation rather than organizational choice; (c) work whose task decomposability is too low for task-level analysis (the proposed threshold is fewer than approximately ten distinct O*NET tasks, or independent-coder disagreement above 30%). Senior executive roles and certain forms of integrated creative direction fall into this last category in practice.

Furthermore, the OOA framework assists in redesigning workforce architecture, it does not forecast how much AI will reshape the economy. The framework adds specificity, and the specifications have two parts. The first is a redistribution prescription for the workforce architecture thru the typology of three archetypes. The second is a pipeline-preservation prescription in order to have line of sight to key skills needed to support the future of work.

This is also what separates the framework from the other approaches now in circulation. The closest practitioner frameworks, including McKinsey's Agentic Organization, Deloitte's Humans × Machines, and Daugherty and Wilson's Human + Machine, engage one of the two prescriptions, or part of one, but none engages both at the level of architectural specificity attempted here. Appendix D outlines the framework in further details across other roles.

2.3 Legislative Challenges, Policy Change and Organizational Imperatives

2.3.1 Can Education and Reskilling Help?

Policymakers and corporations frequently frame workforce preparation strategies and tactics as a skills deficit problem, solvable through training programs, certifications, or structured online courses (US Department of Labor). Reskilling can support some aspects of this transition, but its efficacy will be minimal. For example, certain roles require tacit learning in order to be effective, as formal instruction lacks experiential learning. A training program can teach the vocabulary of credit risk analysis, but it

cannot replicate the accumulated exposure that comes from handling hundreds of cases across different market conditions, under the supervision of experienced practitioners who can identify when a heuristic is being applied incorrectly, or correctly. When AI eliminates the experience-building tasks that historically produced that exposure, tacit learning is removed, and training programs struggle to compensate for the absence of real-world developmental context.

The reskilling optimism also tends to underestimate the scale of the challenge relative to the pace of change. The World Economic Forum (2023) projects that 60 percent of the global workforce will require significant reskilling within the next five years, a figure of institutional capacity for workforce preparation that has never existed. Historical reskilling programs, even well-funded ones, have demonstrated limited effectiveness at scale, and ineffective in environments that move at high speeds.

This thesis thus argues, to have a successful workforce transition in the age of automation, augmentation, and agentic AI, intentional redesign and tacit learning must be a part of the experience developing path in order to develop AI fluency, and co-exist alongside AI.

2.3.2 Why Markets Will Not Self-Correct By Themselves

The organizational redesign described in the previous sections, and any future educational reforms have a common limitation: they depend on individual institutions making conscious choices to prepare their workforce. This can be costly and time-consuming, negating any performance benefits extracted from AI. In short, there is no incentive to do so. That's why favorable policies need to be instituted and pursued with urgency, otherwise we will experience three (3) distinct market failures that, when occurring simultaneously, will create compounding impacts to the disruption already under way.

First is the efficiency urgency documented in Part A. Individual organizations respond to competitive pressure to improve efficiency, increase EPS (earnings per share), and eliminate headcount where work is repetitive and can be automated. This is directly impacting college graduates and early-career adults compressing the workforce. Additionally, the apprenticeship pathways that develop workforce experts is seeing early signs of erosion. No single firm is sufficiently incentivized to pursue intentful redesign and sustain the development pipelines that competitors are dismantling. As such, policies need to be instituted in order to give companies a favorable cost-benefit analysis to justify the investment in the workforce (Muehlemann, Samuel & Wolter, Stefan C., 2011).

Second is a structural unemployment risk. Prior waves of automation displaced workers from tasks while creating new categories of work that absorbed them, typically within a timeframe that allowed institutional adaptation. The evidence reviewed in Part A suggests AI could compress that adaptation window in ways that may produce a sustained net job deficit period, which we will not recover from, particularly for younger and lower-credentialed workers entering the labor market during the transition.

Third firms making AI displacement decisions now, will not bear the full cost of those decisions until well past the tenureship of the executive who made the decision, because current corporate incentive programs support short-term performance gains to executives, not long term. As such, the cognitive development costs identified in this paper will have a lag of five to eleven years, when mastery of mid-level and early-career talent is expected to develop, but never does due to the displacement approaches and decision made today. This is perpetuated by the current market of CEOs for Fortune 1500's where the average performance outlook is documented as three years (Spencer Stuart, 2025). This means, CEOs are incentivized to financially optimize a company in a three-year horizon, and not past that.

These three market failures operate simultaneously and compound each other, and in conditions where durable market self-correction do not exist, and the timeframe for voluntary institutional response has already passed, we're in need of rapid policy change to absorb further consequences.

2.3.3 The Legislative Challenge: Speed, Complexity, and Competing Interests

The case for policy intervention is clear, but the path forwards is complicated. We have found three challenges as to why AI legislation has lagged behind AI deployment, and why catching up will require more than political will.

The first is legislative legal debt. We may be familiar with the concept of technical debt in enterprise; a historic lack of technical updates creates a much sharper update curve, presenting challenges to improving base technology and complicating transformation efforts. Similarly, we are experiencing legal debt, whereby foundational legal infrastructure such as data privacy frameworks, algorithmic accountability standards, employment classification rules are lacking.

The second challenge is the preemption problem. Law is structurally reactive: it codifies responses to harms that have already manifested rather than preventing harms that have not yet occurred. Effective AI governance however requires preemptive action, and preemption is what legislative processes are least equipped to produce.

The third challenge is competition of interests. Any AI policy proposal becomes a contested political object rather than a technical consensus. Organized labor lobbies displacement protections and union authority in AI-related workforce decisions; technology firms lobby for deployment velocity and resists liability frameworks; capital markets lobby the productivity gains that make AI investment economically attractive; and smaller enterprises face adoption costs and competitive pressures that large firms do not. These interests do not resolve cleanly along partisan lines, and partisan dynamics further complicate coalition-building. Any durable AI workforce policy will require managing these tensions explicitly rather than designing around them.

The fourth challenge is geopolitical competition. Unlike previous technological revolutions, artificial intelligence is widely viewed not only as an economic asset but also as a strategic national capability. Policymakers increasingly frame AI through the lens of national defense competitiveness, economic security, and military advantage, creating pressure to accelerate deployment rather than slow it for deliberation. The result is a persistent fear that stronger oversight, transparency requirements, or workforce protections could place one nation at a disadvantage relative to competitors such as China, who are simultaneously investing heavily in AI infrastructure, talent, and deployment capabilities. This dynamic creates what might be described as a *strategic acceleration dilemma*. The perceived costs of moving too slowly often appear more immediate than the risks of moving too quickly. As a result, policymakers, enterprises, and investors may become willing to tolerate higher levels of uncertainty, workforce disruption, and societal risk in exchange for maintaining competitive position.

2.3.4 Eight (8) Frameworks and Policy Changes

As we've seen, the labor-policy challenge created by AI extends beyond the direct worker displacement. Traditional labor shocks typically appear as layoffs, plant closures, or sectoral contractions. By contrast, AI-related workforce restructuring emerges more gradually through reduced entry-level hiring, non-replacement of knowledge workers, workforce pyramid compression, and the transfer of early-career operational work to AI systems and software agents. If AI increasingly absorbs the tasks through which workers historically develop expertise, judgment, and tacit knowledge, then labor policy cannot focus solely on compensating displaced workers after disruption occurs. It must also address the preservation of

capability formation pathways through which future workers acquire the skills necessary to supervise, govern, and work alongside AI systems.

Therefore, this thesis proposes eight (8) frameworks and policy changes to support the labor-market AI transitions, with the objective of not slowing technological progress, but to better align the economic incentives of AI adoption with the long-term sustainability of workforce development and institutional capability formation.

The AI Nutrition Label: Framework for Data Transparency

Prior to 1990, nutritional disclosures on packaged foods in the United States were inconsistent, voluntary, and often opaque to the average consumer. Congress sought to correct this which granted the U.S. Food and Drug Administration (FDA) authority to mandate standardized nutrition labeling and regulate nutrient and health claims on packaged foods. The underlying rationale was both economic and societal: consumers could not make informed decisions about what they were consuming without visibility into the ingredients, composition, and health implications of the products entering their bodies. The Nutrition Labeling and Education Act of 1990 (NLEA) therefore established transparency as a prerequisite for trust, informed choice, and public accountability.

This thesis proposes extending a comparable regulatory framework to Large Language Model and frontier AI providers. As these systems become embedded in decision-making across communication, education, healthcare, finance, hiring, and public administration, society is increasingly asked to trust models whose underlying data remains largely undisclosed. To improve transparency, this thesis advocates for standardized '*AI Data Nutrition Labels'* requiring model providers to disclose key characteristics of their training data, including data sources and provenance, exclusion criteria, the use of real versus synthetic data, known representation gaps, observed biases, hallucination tendencies, model degradation over time, and the conditions under which reliability declines.

Importantly, this framework would not require disclosure of proprietary model weights, trade secrets, or other protected intellectual property. Instead, it would establish a standardized transparency protocol focused on the information necessary for public trust and institutional accountability. Required disclosures could include training-data composition by source, geography, language, modality, and time horizon; safety testing results and benchmark limitations; reinforcement-learning and human-feedback methodologies; copyright, consent, and data-provenance practices; computational and environmental impacts; and model versioning and update histories.

Ultimately, the AI Nutrition Label reinforces a foundational principle that innovation and trust are not competing objectives, but mutually reinforcing conditions for sustainable technological progress.

AI Agent Payroll Tax

Once AI deployment crosses a defined workforce impact threshold, firms should be required to make an AI agent payroll tax. The proposal should not be understood as a generalized tax on software, AI models, or innovation itself. Rather, it functions as a payroll-equivalent adjustment that applies when AI systems materially substitute human labor and the corresponding employee labor tax. The policy rationale follows an established principle within existing labor-market systems: employers already contribute financially to public structures designed to absorb labor-market risk. Employers also pay federal unemployment taxes

under FUTA, which, together with state unemployment systems, supports unemployment compensation for workers who lose their jobs (Internal Revenue Service, 2026b). When firms hire human labor, they pay in an employee tax, and that tax contributes to a fiscal structure that supports retirement security, health insurance for older adults, and unemployment protection. Therefore, this thesis proposes, under current policy structures, firms also incur payroll-related obligations when AI labor systems replace or absorb human labor systems. Framing the proposal as a payroll-equivalent contribution: partially restoring neutrality between human labor and automated labor in cases where AI directly substitutes for workers. Additionally, the revenue generated through this contribution mechanism could support a variety of other causes. It could finance direct workforce transition programs, including retraining, relocation assistance, temporary wage support, and wage insurance for displaced workers. It could support apprenticeships, paid internships, rotational programs, fellowships, and early-career AI-supervision pathways designed to preserve capability formation within AI-exposed occupations. A portion of the fund could support human-centered sectors such as healthcare, education, mental health, caregiving, and community services, where human presence, empathy, relational continuity, and trust remain central sources of value creation. The AI agent payroll tax should be understood as a stabilization mechanism rather than an anti-technology measure.

Early-Career Hiring Tax Credits: Developing the Experts of the Future

The proposed AI agent payroll contribution should be paired with a complementary positive incentive: a Tax Credit for those who continue to hire and train early-career adults. While payroll tax addresses the downstream costs associated with AI-driven labor substitution, the tax credit addresses the developmental pathways needed to form expertise in early-career adults. As firms face increasing competitive pressure to automate operational and entry-level work, even organizations that recognize the long-term value of developing junior talent may reduce those roles to remain cost competitive. The result is a classic collective-action problem. Individual firms will benefit in the short term by reducing developmental labor costs, while the broader economy loses the institutional pathways through which expertise, judgment, and future supervisory capability are formed. The use of tax incentives to align private incentives with broader public objectives is already well established within U.S. economic policy. As such, our proposal extends what is in place so we have infrastructure support.

Existing examples include the Work Opportunity Tax Credit, which incentivizes hiring among populations facing employment barriers. The Work Opportunity Tax Credit gives employers a federal tax credit for hiring individuals from targeted groups that face barriers to employment (Internal Revenue Service, n.d.). The Employee Retention Credit was designed during the COVID-19 crisis to encourage eligible employers to keep workers on payroll despite economic hardship (Internal Revenue Service, 2025). Registered Apprenticeship combines paid on-the-job training with classroom instruction (U.S. Department of Labor, n.d.-c). Clean-energy tax credits also show how Congress has attached labor standards, including prevailing wage and apprenticeship requirements, to broader industrial-policy goals (U.S. Department of Energy, n.d.). Although these programs were not designed specifically for AI transition, they establish an important policy principle: the tax system can be used to encourage forms of employment that markets may otherwise underproduce despite their broader societal value.

The objective of the proposed tax credit would support roles that preserve the developmental functions most vulnerable to AI compression, including structured entry-level programs, apprenticeships, paid internships, rotational assignments, trainee cohorts, and supervised AI-assisted roles where workers learn not only to use AI systems, but also to evaluate, challenge, and govern them. Importantly, eligibility should depend on the developmental quality of the role rather than simple headcount preservation. To qualify, organizations would need to demonstrate that roles include structured supervision, meaningful exposure to substantive work, responsible AI training, and a credible pathway toward higher-skill capability development. A junior employee reduced primarily to passively validating AI outputs without

understanding the underlying work would not meet the standard. By contrast, roles involving research, drafting, analysis, client interaction, AI-assisted production, and supervised human review would. The policy objective is not preserving employment for its own sake but preserving the capability formation processes through which future expertise is developed.

The credit would need guardrails. It should be targeted to AI-exposed occupations where entry-level hiring has declined and the occupation is central to future AI oversight. It should be refundable or transferable for small and medium-sized firms, which may face stronger cost pressure than large firms. It should include an anti-substitution rule so that firms cannot claim the credit for creating a small number of eligible roles while eliminating a larger number of similar developmental roles through AI deployment. It should also require reporting on roles, occupations, training provided, AI systems used, retention rates, and progression outcomes. The credit should also be linked to the AI Workforce Impact Threshold. Firms that preserve or expand entry-level roles after AI deployment could receive favorable treatment. Firms that cross the threshold for labor substitution without maintaining developmental roles would face the contribution. This creates a balanced structure: policy does not tell firms to avoid AI, but it shifts the payoff structure around how AI is implemented. Firms can adopt AI and receive support if they use it to develop workers. Firms that use AI primarily to remove workers contribute to the transition system.

The case for the credit is strengthened by evidence that early labor-market shocks can have long-lasting consequences. Oreopoulos, von Wachter, and Heisz find that graduates entering the labor market during a recession experience substantial and unequal earnings losses, with persistent effects lasting many years, particularly for lower-earning graduates (Oreopoulos et al., 2012). This evidence is not about AI specifically, but the analogy is important. If AI suppresses entry-level opportunities across multiple graduating cohorts, the consequences may persist beyond the initial period of technological adoption.

This is the most direct policy connection to the paper's central thesis. AI may remove not only jobs, but the work through which people become capable of doing harder jobs. If the market undervalues developmental work because AI can produce the immediate output more cheaply, public policy can help price in the long-term social value of that work. The credit turns an abstract concern about talent pipelines into an operational fiscal mechanism.

It's important to note, the thesis does not encourage preserving outdated tasks nostalgically simply because they existed historically. Many operational activities will and should become transformed. The policy objective instead is to deliberately rebuild the developmental function to acquire the expertise necessary to supervise, govern, and collaborate effectively with increasingly intelligent systems.

AI Deployment Registries

One of the defining governance gaps of the AI era is that society currently lacks a standardized mechanism for tracking frontier AI systems once they move from proof of concept into production, even though governments maintain oversight frameworks for pharmaceuticals, financial institutions, environmental disclosures, and securities markets. No equivalent infrastructure exists for monitoring high-impact AI deployments despite their growing influence on labor markets, economic systems, public services, and daily human decision-making. The absence of visibility creates a significant policy blind spot.

To address this gap, this thesis proposes the creation of a federally mandated AI Deployment Registry for enterprise AI systems that exceed defined thresholds of capability, workforce impact, operational autonomy, or societal influence. Organizations deploying qualifying systems would be required to register their systems when their deployments materially alter workforce composition, automate economically significant workflows, influence public-facing decisions, or operate within critical sectors such as healthcare, finance, education, infrastructure, legal services, defense, and public administration.

The purpose of the registry is not to restrict innovation, but to establish visibility. Similar to the role SEC corporate filings play in capital markets, or FDA's systems for reporting adverse events, medication errors, and product problems play via their FDA Adverse Event Monitoring System (AEMS/FAERS), the registry would create a transparent record of high-impact AI deployments, enabling policymakers, researchers, economists, and the public to better understand how artificial intelligence is reshaping industries, occupations, and economic activity in real time. Registry data would be reviewed by an independent multidisciplinary body tasked with assessing aggregate trends in workforce displacement, wage compression, occupational restructuring, productivity concentration, and emerging capability risks. Over time, this would create one of the first national datasets capable of measuring AI's macroeconomic effects using observed deployment data rather than forecasts, vendor-sponsored studies, or anecdotal evidence. Such visibility would strengthen evidence-based policymaking and provide governments with the institutional capacity to respond to emerging risks before they become systemic.

Additionally, drawing inspiration from the United Kingdom's Financial Conduct Authority regulatory sandbox model for fintech innovation, this proposal also suggests a parallel "AI Sandbox Pathway" for organizations seeking accelerated deployment approval. Under this pathway, companies could voluntarily partner with the federal oversight body during development and testing phases within a monitored sandbox environment. By integrating transparency, workforce-impact assessments, and safety evaluation earlier in the deployment lifecycle, participating organizations could receive accelerated approval while maintaining public accountability.

Together, the AI Deployment Registry and Sandbox Pathway create a framework that balances and promotes the velocity of innovation with visibility.

AI-ISO Standards

As artificial intelligence becomes increasingly embedded in hiring, workforce management, performance evaluation, and organizational planning, a critical governance gap has emerged in that, there is no universally recognized operational standard governing how AI systems should be designed, deployed, audited, or monitored when they materially influence employment outcomes. While industries ranging from manufacturing to cybersecurity rely on internationally recognized standards to establish requirements for quality, safety, accountability, and interoperability, workforce-related AI systems remain governed largely by fragmented internal policies, inconsistent vendor practices, and voluntary ethical guidelines. Yet these systems increasingly influence who gets hired, promoted, retrained, managed, or displaced.

To address this gap, this thesis proposes the formation of the AI Workforce Deployment Standard (AWDS), an international framework modeled after International Organization for Standards (ISO), yet designed specifically for workforce-impacting AI systems. The standard would apply to AI deployments used in hiring, recruiting, workforce planning, productivity monitoring, performance management, compensation modeling, organizational restructuring, and large-scale workforce reductions. Its purpose is not to restrict innovation, but to establish a common set of principles, controls, and accountability mechanisms governing how AI is introduced into employment-related environments. A central requirement of AWDS would be the completion of a Workforce Impact Assessment (WIA) for any AI deployment expected to materially affect employment outcomes or workforce composition above a certain threshold, we recommend 75 or 1% of total workforce to be that threshold. Prior to deployment, organizations would be required to disclose anticipated workforce impacts and required to perform a 'role impact analysis' and illustrates the affected functions, departments, roles, and tasks. They would also include a 'Workforce Transition Plan' that highlights retraining, transition, and worker supported programs, along with a 'Severance and Support Plan' that outlines the severance frameworks, career

transition assistance, educational stipends and placement resources, which currently exist today for individuals laid off. The objective is to ensure that workforce transformation is managed transparently rather than occurring through opaque operational decisions justified solely by efficiency gains.

Extend Antitrust Laws to Private AI companies

Modern antitrust laws were created to prevent excessive concentrations of economic power from distorting markets, suppressing competition, and creating societal harm. The Sherman Antitrust Act of 1890 and subsequent antitrust frameworks, like the Clayton Antitrust Act of 1914, were designed to preserve competitive markets, reduce systemic dependencies, and protect the public interest from concentrated private control. However, these frameworks were developed for an industrial economy and were not designed to address the unique forms of concentration emerging in the age of artificial intelligence.

Unlike previous monopolistic structures built around physical infrastructure or distribution networks, frontier AI firms increasingly control the foundational cognitive infrastructure upon which future economies may depend. A small number of organizations now possess disproportionate influence over advanced models, training data, compute resources, talent, and the AI platforms increasingly embedded within enterprise software, education, healthcare, finance, government, and public services. Yet many of the most influential firms remain privately controlled entities with limited disclosure obligations despite possessing societal influence comparable to major public infrastructure providers.

Thus, this thesis proposes the modernization of antitrust frameworks to include large-scale private AI firms whose market concentration exceeds defined thresholds of model capability, compute ownership, enterprise dependency, or ecosystem influence. Organizations meeting these thresholds would be designated as Systemically Significant AI Entities (SSAIEs) and become subject to enhanced transparency, market-concentration reviews, and public-interest accountability requirements regardless of ownership structure. The rationale is to avoid excessive concentration of AI capability and encourage competition, increase innovation diversity, strengthen startup viability, and increase systemic dependence on a small number of providers. Under this framework, qualifying firms would be required to provide periodic disclosures regarding model deployment scale, compute concentration, infrastructure partnerships, enterprise dependency metrics, workforce impacts, safety incidents, acquisitions, and major commercial agreements. Regulators would additionally possess authority to review mergers, exclusivity arrangements, and other practices capable of undermining competition within the AI ecosystem.

The purpose of this policy is to ensure that the future of AI remains competitive, resilient, and aligned with the broader public interest rather than concentrated within a small number of private actors.

Displacement Notification Requirements

One of the most direct policy interventions involves extending existing workforce notification frameworks to explicitly address AI-driven restructuring. Current labor regulations (U.S. Department of Labor, n.d.) generally require advance notice for large-scale layoffs but do not distinguish between cyclical workforce reductions and structural changes resulting from automation or artificial intelligence. As a result, many forms of AI-related displacement remain largely invisible until their effects are already embedded within the labor market. The most relevant precedent is the Worker Adjustment and Retraining Notification Act (WARN), which requires employers to provide advance notice before plant closures or mass layoffs.

We propose a similar act for AI-induced workforce adjustments. Because right now AI-related workforce restructuring often occurs with direct job loss, but mostly through less visible mechanisms that include hiring suppression, workforce pyramid compression, reducing internship and apprenticeship opportunities, without ever triggering a conventional layoff event. In following the principles outlined in WARN, our suggested policy would require an 'Impact Report' that discloses when role reductions or workforce compression occurs in a concentrated area of the company. The disclosure would categorize AI systems deployed and the functions affected, the career levels exposed, the number of roles eliminated and replaced. Furthermore, the report would include where reductions disproportionately affect junior or developmental roles.

Under this framework, organizations would be subject to disclosure requirements based on the scale and nature of workforce impact. Initial disclosures would provide visibility into AI-related labor substitution, hiring reductions, affected occupations, and restructuring events, to provide insights into formal workforce-impact assessments evaluating retraining, redeployment, and workforce-transition alternatives.

Autonomous Capability Threshold (Kill Switch)

Historically, governments have imposed heightened oversight on technologies capable of producing systemic consequences beyond the control of their operators. Nuclear facilities, aviation systems, pharmaceuticals, and consumer products, are all subject to additional scrutiny when their capabilities exceed established safety thresholds. Artificial intelligence presents a similar challenge.

This thesis therefore proposes the establishment of an Autonomous Capability Threshold (ACTF), under which AI systems would be subject to mandatory review, operational restrictions, or temporary suspension when they demonstrate predefined capabilities beyond established safety thresholds. These capabilities include the ability to substantially improve their own performance without human authorization, coupled with establish or pursue self-derived goals, execute consequential actions autonomously, or acquire computational, financial, informational, or physical resources necessary to expand their capabilities. The purpose of this threshold is not to constrain innovation, but to create a capability-based oversight mechanism that activates when AI systems begin exhibiting characteristics associated with reduced human control. By establishing clear intervention thresholds, we autonomous AI systems remain subject to meaningful oversight, accountability, and alignment with human intent before they become deeply embedded within critical societal infrastructure.

Policy Change Summary

Collectively, these eight (8) policy proposals represent an integrated governance framework for managing the transition into an AI-driven economy. Rather than treating artificial intelligence solely as a technological advancement, the framework recognizes AI as a systemic force, reshaping labor market, institutional structures, economic participation, and the cognitive infrastructure of society itself. The proposals therefore seek to modernize the relationship between innovation, accountability, workforce development, and public policy before existing systems become structurally obsolete.

First, the thesis proposes the establishment of an AI Nutrition Label Framework requiring frontier AI developers to disclose standardized information regarding training data composition, synthetic versus real data usage, known biases, safety limitations, drift patterns, and model reliability constraints. Inspired by the Nutrition Labeling and Education Act of 1990, the proposal reframes transparency as a prerequisite for public trust rather than an obstacle to innovation.

Second, the thesis introduces an AI Agent Payroll Contribution**,** designed as a payroll-equivalent adjustment mechanism for organizations materially substituting human labor with AI labor. Rather than functioning as a generalized tax on software or innovation, the mechanism ensures that organizations benefiting from large-scale automation contribute proportionally to "social insurance" programs, similar to employee tax, that fund and supports Social Security, Medicare, and unemployment insurance, social programs that provide safety nets for workers.

Third, the thesis proposes Early-Career Hiring and Human Capital Formation Tax Credits to incentivize firms to preserve apprenticeships, entry-level pathways, rotational programs, and supervised developmental work. This policy directly addresses one of the central findings of the thesis that AI threatens not only jobs, but the capability formation processes through which expertise, judgment, and future supervisory capacity are developed.

Fourth, the thesis advances the creation of AI Deployment Registries**,** requiring high-impact AI systems exceeding defined capability or workforce-impact thresholds to be registered with a federal authority. These registries would establish public visibility into large-scale AI deployments while creating one of the first institutional infrastructures capable of monitoring AI's macroeconomic, labor-market, and societal effects in real time. Complementing this structure is a regulatory sandbox pathway modeled after U.K. fintech regulation, allowing organizations to collaborate proactively with regulators in monitored development environments.

Fifth, the thesis proposes an AI Workforce Deployment Standard (AWDS) governed through an AI-specific ISO-equivalent body. This framework would establish internationally recognized operational standards governing the use of AI in hiring, workforce planning, performance management, restructuring, and labor displacement decisions. Embedded within this proposal is the requirement for a mandatory Workforce Impact Assessment (WIA), ensuring that organizations deploying workforce-impacting AI systems publicly disclose role impacts, retraining plans, transition support structures, worker participation mechanisms, and post-deployment review outcomes.

Sixth, the thesis recommends extending Antitrust Oversight to Systemically Significant Private AI Firms**.** Existing antitrust frameworks were designed primarily around public industrial monopolies; however, frontier AI firms increasingly control the cognitive infrastructure upon which modern economies depend. This proposal introduces enhanced transparency and oversight requirements for dominant private AI entities whose concentration of compute, model capability, infrastructure, and market influence may create systemic economic, social, or geopolitical risks.

Seventh, the thesis proposes modernizing Displacement Notification Requirements through AI-specific extensions of existing workforce disclosure frameworks such as WARN. Unlike traditional layoffs, AI-related labor restructuring often occurs incrementally through hiring suppression, role compression, and non-replacement of early-career workers. The proposal therefore introduces workforce-impact thresholds, pipeline triggers, Workforce Impact Notices, and formal AI Workforce Impact Assessments capable of identifying structural capability erosion before it becomes economically irreversible.

Eight, the thesis proposes a *kill switch* for any autonomous AI systems that demonstrate self-evolving capabilities such as setting its own goals, acquiring its own resources, improving its own code base, and acting autonomously and independently without human governance.

Together, these eight constitute a governance architecture for a transition that is already underway and a framework that treats AI not as a productivity tool to be deployed, but as a structural force to be governed since the age of automation, augmentation and agentic AI demands institutional frameworks capable of

governing not only productivity gains, but also human transition, workforce resilience, economic stability, and societal trust.

2.4 Preserving Human Capability in the AI Economy

2.4.1 What the Evidence Demands

This paper set out to answer a question that is deceptively simple in its formulation and genuinely difficult in its implications: what happens to human work, human development, and human capital as we pivot into the era of automation, augmentation, and agentic AI? The analysis produced seven findings that, taken together, constitute a structural diagnosis.

First, AI-driven disruption is asymmetric, not universal. The evidence supports neither catastrophic mass unemployment nor broadly shared technological abundance. Instead, it illustrates that AI concentrates disruption on routine, junior, and experience-building work while amplifying the productivity of workers who have already accumulated the competence. The result is not immediate mass displacement but structural compression and gradual elimination of the developmental pathways through which labor markets renew themselves.

Second, historical recovery mechanisms are impaired. Prior disruptions allowed institutions to absorb displacement because new industries emerged, new roles were created, and training systems had time to respond. AI compresses this sequence on two fronts simultaneously: advancing displacements faster than institutions can track, while automating the experiential work through which people have historically built the competence to take on harder roles. The pathway from displacement to stabilization is now narrower and less forgiving than historical analogies suggest.

Third, the deepest risks are still invisible. AI will absorb routine work, generate productivity gains, and reduce unit costs. The consequences however aren't fully understood as trends around cognitive offloading, apprenticeship pipeline erosion, and the substitution of nominal oversight for genuine judgment erode. These are the deferred costs of successful deployments that will not be visible for many years.

Fourth, roles are the unit of transformation. The future of work is will not be determined by which titles survive, but by which human contributions remain irreplaceable within AI-reshaped workflows. The Orchestrator, Operator, and Artisan archetypes represent three durable modes of human value creation, each requiring distinct capabilities, governance structures, and developmental pathways. Organizations that redesign around them are building for what the post-AI economy requires.

Fifth, reskilling is insufficient. Workforce strategies focused on tools and certifications do not address the skills mismatch associated with future requirements of roles. Workforce preparation is a systems-design challenge and requires role redesign, apprenticeship reconstruction, and the deliberate embedding of developmental exposure inside AI-mediated workflows.

Sixth, compression of the workforce pyramid and the pathway for expertise development. This is the paper's most consequential finding. AI can amplify inequality, compress career pathways, and erode the cognitive development by which sound judgment is formed if policy changes, educational institutions and governance protocols aren't instituted before the structural patterns forming in labor markets and organizational practice harden into conditions that are difficult to reverse.

Seventh, the talent formation fracture is a defining risk in the labor market. To sustain a pipeline of experts capable of orchestrating, operating, and governing AI systems, organizations must continue hiring and investing in early-career talent. Early-career workers denied entry-point lose the developmental pathway through which judgment, intuition, and oversight capacities are built.

2.4.2 The Workforce 4.0 Imperative

Workforce 4.0 is defined as the deliberate redesign of how humans learn, develop, and contribute alongside AI systems across organizations, educational institutions, and policy environments to sustain the formation of human expertise, judgment, and oversight capacity in an era of automation, augmentation, and agentic AI. Workforce 4.0 considers what must be deliberately protected, built, and governed and what structural scaffolding is required to ensure that human capability compounds rather than erodes as AI scales.

The frameworks suggested emerged from the convergence of four analytical streams documented: labor market evidence showing asymmetric disruption concentrated at the base of the workforce pyramid; geographic and demographic data revealing uneven adaptive capacity across economies and cohorts; cognitive research demonstrating that AI-assisted work reduces the neural engagement through which durable expertise is formed; and organizational evidence showing that workflow redesign, not deployment alone, determines whether AI augments or substitutes for human judgment over time. Workforce 4.0 is provides an evolving blueprint designed to help institutions detect misalignment earlier and intervene before disruption hardens into irreversible loss.

The framework rests on four structural pillars. Each is a prerequisite and none are sufficient alone.

Pillar I: Organizational Redesign. Firms must treat AI deployment as a workforce re-architecture decision, not a technical deployment. This requires applying the OOA role framework to redistribute work in ways that preserve and amplify human capital in order for companies to extract the most value. It requires the deliberate reservation of developmental work for early-career workers within AI-mediated workflows. It requires discernment to understand not all entry-level tasks should be automated simply because they can be because some tasks build the intuition, pattern recognition, and contextual judgment that organizations will require at senior levels within five to eleven years.

Pillar II: Talent Development Architecture. The M-shaped talent taxonomy, which is a composition of domain expertise, systems thinking and workflow design, AI literacy, and human judgment, defines the capability profile required of workers in AI-augmented environments. Organizations must rebuild apprenticeship structures, rotational programs, and supervised early-career roles that expose workers to consequential decisions under conditions of genuine uncertainty which classroom instruction and certification programs cannot replicate.

Pillar III: Measurement. This thesis proposes a Human Capital Health Index, a composite institutional metric tracking early-career hiring rates, apprenticeship pathway preservation, cognitive task retention in human roles, employee progression from entry to senior levels, and AI override frequency as a proxy for substantive versus nominal human oversight. What gets measured gets managed. Without a metric, capability erosion remains invisible until it materializes as a governance failure.

Pillar IV: Policy Alignment. Organizational redesign and talent architecture alone are insufficient against the incentive structure that governs most AI deployment decisions. Individual firms will rationally thin junior pipelines while the collective cost, a generation of workers denied the developmental rung through which expertise is built, accumulates at the labor market level. The seven policy proposals in Section 2.4

constitute the external scaffolding that realigns those incentives: payroll contributions that price in the transition cost of labor substitution, tax credits that reward developmental hiring, deployment registries that create institutional visibility, and workforce impact standards that impose accountability before restructuring occurs rather than after.

Every major technological transition has tested humanity's capacity to redesign its institutions before the dislocations compound into irreversible harm. The industrial revolution eventually produced labor law, public education, and social insurance, but largely after the human cost had already been paid. The digital revolution's institutional response remains contested and incomplete. AI presents the same test at greater speed and with less tolerance for delay. It is faster, more cognitively intimate, and more structurally penetrating than any prior general-purpose technology because it does not merely change what work is done, it changes the cognitive processes through which people become capable of doing harder work.

The workforce challenge AI poses is about whether the institutions through which human develop competence, build identity, and contribute to collective enterprise will be redesigned to remain functional in an environment where the work that historically performed those functions is being automated at scale. While we debate if intelligence is scalable, meaning, judgment, and responsibility are not. The future of work will be decided less by what AI can do than by what institutions choose to preserve, what organizations choose to measure, and what policymakers choose to protect as the transition unfolds. Those choices, not the technology itself, are the defining challenge of the current era, and they demand a response now, before the structural conditions forming in labor markets, educational systems, and organizational practice have compounded into constraints that cannot be reversed.

References

A

Accenture. (2023, January 17). Accenture identifies emerging group of industry leaders adopting 'Total Enterprise Reinvention' as a strategy to reach new performance frontier [Press release].https://newsroom.accenture.com/news/2023/accenture-identifies-emerging-group-of-industry-leaders-adopting-total-enterprise-reinvention-as-a-strategy-to-reach-new-performance-frontier

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487. National Bureau of Economic Research. https://www.nber.org/papers/w32487

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3-30. https://doi.org/10.1257/jep.33.2.3

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. Journal of Labor Economics, 40(S1), S293–S340. https://doi.org/10.1086/718327

AI Safety Summit. (2023, November 1). The Bletchley Declaration by countries attending the AI Safety Summit, 1–2 November 2023. UK Government. https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. Labour Economics, 71, Article 102002. https://doi.org/10.1016/j.labeco.2021.102002

Allianz Trade. (n.d.). Retail in the US: Towards destructive destruction. Allianz Trade.https://www.allianz-trade.com/en_global/news-insights/economic-insights/retail-in-the-us-towards-destructive-destruction.html

American Medical Association. (2025). Doctors work fewer hours, but the EHR still follows them home. AMA Organizational Biopsy 2024 National Report. https://www.ama-assn.org/practice-management/physician-health/doctors-work-fewer-hours-ehr-still-follows-them-home

Amiti, M., Redding, S. J., & Weinstein, D. E. (2019). The impact of the 2018 tariffs on prices and welfare. Journal of Economic Perspectives, 33(4), 187–210. https://doi.org/10.1257/jep.33.4.187

Andrade, Philippe, Alexander Dietrich, Sophie Handley, John Leer, Raphael Schoenle, Jenny Tang, and Egon Zakrajšek. 2025. “Effects of Tariff Uncertainty on the Outlook of Small and Medium-sized Businesses.” Federal Reserve Bank of Boston Current Policy Perspectives 25-12. Published September 5, 2025. https://www.bostonfed.org/publications/current-policy-perspectives/2025/tariff-uncertainty-on-small-and-medium-businesses.aspx

Anthony, C., Bechky, B. A., & Fayard, A.-L. (2023). "Collaborating" with AI: Taking a system view to explore the future of work. Organization Science, 34(5), 1672–1694. https://doi.org/10.1287/orsc.2022.1651

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. https://doi.org/10.1257/jep.29.3.3

Autor, D. H. (2024). Applying AI to rebuild middle class jobs (NBER Working Paper No. 32140). National Bureau of Economic Research. https://doi.org/10.3386/w32140

Autor, D. H., & Thompson, N. (2025). Expertise. NBER Working Paper No. 33941. National Bureau of Economic Research. https://www.nber.org/papers/w33941

B

Bailey, D. E., & Barley, S. R. (2020). Beyond design and use: How scholars should study intelligent technologies. Information and Organization, 30(2), Article 100286. https://doi.org/10.1016/j.infoandorg.2019.100286

Bangladesh Labour Foundation (BLF), BRAC University & Solidaridad Network Asia (2025). Assessment of Technological Transition in the Apparel Sector of Bangladesh and Its Impact on Workers. Automation caused a 31% decline in the RMG sector workforce; cutting stage saw a 48% workforce reduction; sweater factories 37%.https://www.just-style.com/news/bangladesh-rmg-workers-automation/

Bank of England. (n.d.). Research datasets. Bank of England Statistics.https://www.bankofengland.co.uk/statistics/research-datasets

Beane, M. (2019). Shadow learning: Building robotic surgical skill when approved means fail. Administrative Science Quarterly, 64(1), 87–123. https://doi.org/10.1177/0001839217751692

Bedard, J., Kropp, M., Hsu, M., Karaman, O. T., Hawes, J., & Rosen Kellerman, G. (2026, March 5). When using AI leads to "brain fry." Harvard Business Review.https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry

Betsy Sparrow et al. ,Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips.Science333,776-778(2011).DOI:10.1126/science.1207745

Bidwell, M. (2011). Paying more to get less: The effects of external hiring versus internal mobility. Administrative Science Quarterly, 56(3), 369–407.https://faculty.wharton.upenn.edu/wp-content/uploads/2012/03/Paying_More_ASQ_edits_FINAL.pdf

Board of Governors of the Federal Reserve System. (2011). SR 11-7: Guidance on model risk management.https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm

Bosch, G. (2004). Towards a new standard employment relationship in Western Europe. British Journal of Industrial Relations, 42(4), 617–636.https://doi.org/10.1111/j.1467-8543.2004.00333.x

Boston Consulting Group. (2022, December 7). Five rules for fixing AI in business. BCG.https://www.bcg.com/publications/2022/5-rules-for-fixing-ai-and-machine-learning-for-your-business

Boston Consulting Group. (2024). CEO's guide to maximizing value potential from AI in 2024.https://media-publications.bcg.com/BCG-Executive-Perspectives-CEOs-Guide-to-Maximizing-Value-from-AI-EP0-3July2024.pdf

Boston Consulting Group. (2025). How agents are accelerating the next wave of AI value creation. https://www.bcg.com/publications/2025/agents-accelerate-next-wave-of-ai-value-creation

Boston Consulting Group. (2025, June 26). AI at Work 2025: Momentum builds, but gaps remain.https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain

Boston Consulting Group. (2026). AI will reshape more jobs than it replaces. BCG Henderson Institute. https://www.bcg.com/publications/2026/ai-will-reshape-more-jobs-than-it-replaces

Boyatzis, R. E. (1998). Transforming qualitative information: Thematic analysis and code development. Sage Publications.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77–101. https://doi.org/10.1191/1478088706qp063oa

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab Working Paper.https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/

Brynjolfsson, E., Chandar, B., & Chen, R. (2026, February 9). Canaries, interest rates, and timing: More on the recent drivers of employment changes for young workers. Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/news/canaries-interest-rates-and-timinga-more-on-recent-drivers-of-employment-changes-for-young-workers/

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889–942.https://doi.org/10.1093/qje/qjae044

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? AEA Papers and Proceedings, 108, 43–47. https://doi.org/10.1257/pandp.20181019

Brynjolfsson, E., Rock, D., & Syverson, C. (2018). The productivity J-curve: How intangibles complement general purpose technologies. NBER Working Paper 25148. National Bureau of Economic Research.https://www.nber.org/system/files/working_papers/w25148/w25148.pdf

Bureau of Labor Statistics. (2020). Forty years of falling manufacturing employment. U.S. Department of Labor.https://www.bls.gov/opub/btn/volume-9/forty-years-of-falling-manufacturing-employment.htm

Bureau of Labor Statistics. (2023). Unemployment rate returned to its prepandemic level in 2022. Monthly Labor Review. U.S. Department of Labor.https://www.bls.gov/opub/mlr/2023/article/unemployment-rate-returned-to-its-prepandemic-level-in-2022.htm

Bureau of Labor Statistics. (2023a). Sector 51 - Information - May 2023 OEWS industry-specific occupational employment and wage estimates. U.S. Department of Labor.https://www.bls.gov/oes/2023/may/naics2_51.htm

Bureau of Labor Statistics. (2023b). Sector 52 - Finance and Insurance - May 2023 OEWS industry-specific occupational employment and wage estimates. U.S. Department of Labor.https://www.bls.gov/oes/2023/may/naics2_52.htm

Bureau of Labor Statistics. (2023c). Sector 54 - Professional, Scientific, and Technical Services - May 2023 OEWS industry-specific occupational employment and wage estimates. U.S. Department of Labor.https://www.bls.gov/oes/2023/may/naics2_54.htm

Bureau of Labor Statistics. (2023d). Sector 55 - Management of Companies and Enterprises - May 2023 OEWS industry-specific occupational employment and wage estimates. U.S. Department of Labor.https://www.bls.gov/oes/2023/may/naics2_55.htm

Bureau of Labor Statistics. (2024). Employee tenure in 2024. U.S. Department of Labor.https://www.bls.gov/news.release/pdf/tenure.pdf

Bureau of Labor Statistics. (2025a). Business and financial occupations. Occupational Outlook Handbook. U.S. Department of Labor.https://www.bls.gov/ooh/business-and-financial/

Bureau of Labor Statistics. (2025b). Top executives. Occupational Outlook Handbook. U.S. Department of Labor.https://www.bls.gov/ooh/management/top-executives.htm

Bureau of Labor Statistics. (n.d.-a). Current Population Survey data [Data query tool]. U.S. Department of Labor.https://data.bls.gov/pdq/SurveyOutputServlet

Bureau of Labor Statistics. (various years). Labor force statistics from the Current Population Survey [Data table]. U.S. Department of Labor.https://www.bls.gov/cps/cpsaat01.htm

Burns, T., & Stalker, G. M. (1961). The management of innovation. Tavistock.

C

Catalini, C., Hui, X., & Wu, J. (2026). Some simple economics of AGI. arXiv preprint arXiv:2602.20946.https://arxiv.org/abs/2602.20946

Cevik, S., Fan, A., & Naik, S. (2025). Monetary shocks and labor markets: Evidence from online job vacancy postings (IMF Working Paper No. 2025/058). International Monetary Fund.https://doi.org/10.5089/9798229004251.001

Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). The GenAI divide: State of AI in business 2025. MIT Project NANDA.https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

Colorado General Assembly. (2024). SB 24-205: Consumer protections for artificial intelligence.https://leg.colorado.gov/bills/sb24-205

Consumer Financial Protection Bureau. (2022). CFPB issues guidance on credit denials by lenders using artificial intelligence.https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-on-credit-denials-by-lenders-using-artificial-intelligence/

D

Dani Rodrik, 2016. "Premature deindustrialization," Journal of Economic Growth, vol 21(1), pages 1-33. https://www.nber.org/system/files/working_papers/w20935/w20935.pdf?utm_source=PANTHEON_STRIPPED

Daugherty, P. R., & Wilson, H. J. (2018). Human + machine: Reimagining work in the age of AI. Harvard Business Review Press.

Daugherty, P. R., & Wilson, H. J. (2024). Human + machine, updated and expanded: Reimagining work in the age of AI. Harvard Business Review Press.

David, P. A. (1990). The dynamo and the computer: An historical perspective on the modern productivity paradox. American Economic Review, 80(2), 355–361.https://www.jstor.org/stable/2006600

Davis, J. S. (2026, February 24). AI is simultaneously aiding and replacing workers, wage data suggest. Federal Reserve Bank of Dallas Economic Research.https://www.dallasfed.org/research/economics/2026/0224

De Neve, J.-E., Hancock, J. T., & Niederhoffer, K. (2026, April 15). Why companies that choose AI augmentation over automation may win in the long run. Harvard Business Review.https://hbr.org/2026/04/why-companies-that-choose-ai-augmentation-over-automation-may-win-in-the-long-run

Dekoninck E, Bridge L. THE T-SHAPED DESIGN ENGINEER – USING COHORTS TO EXPLORE HOW SKILLS PROFILES DIFFER THROUGH CAREER STAGES. Proceedings of the Design Society. 2023;3:3533-3542. doi:10.1017/pds.2023.354

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper 24-013.https://www.nber.org/system/files/working_papers/w31161/w31161.pdf

Deloitte. (2024a). The velocity paradox: Chasing AI's potential in the age of speed. Deloitte Insights.https://www2.deloitte.com/us/en/insights/focus/technology-and-the-future-of-work/ai-velocity-paradox.html

Deloitte. (2024b). 2024 Deloitte AI in the boardroom survey. Deloitte.https://www2.deloitte.com/us/en/pages/deloitte-private/articles/ai-in-the-boardroom.html

Deloitte. (2025). AI ROI: The paradox of rising investment and elusive returns. Deloitte Global.https://www.deloitte.com/global/en/issues/generative-ai/ai-roi-the-paradox-of-rising-investment-and-elusive-returns.html

Deloitte. (2025). Work redesign essential to realize AI return on investment [Press release].https://www.deloitte.com/us/en/about/press-room/work-design-essential-to-ai-roi.html

Deming, D. J. (2017). The growing importance of social skills in the labor market. The Quarterly Journal of Economics, 132(4), 1593–1640.https://doi.org/10.1093/qje/qjx022

Demirkan, H., & Spohrer, J. (2015). T-shaped innovators: Identifying the right talent to support service innovation. Research-Technology Management, 58(5), 12–15. https://doi.org/10.5437/08956308X5805007

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances, 10(28), Article eadn5290. https://doi.org/10.1126/sciadv.adn5290

Doty, D. H., & Glick, W. H. (1994). Typologies as a unique form of theory building: Toward improved understanding and modeling. Academy of Management Review, 19(2), 230–251. https://www.jstor.org/stable/258704

DX Research / Laura Tacho (February 2026). Survey of 121,000 developers across 450+ companies. AI-authored code accounts for 26.9% of all production code across 4.2 million tracked developers; 92.6% use AI tools at least monthly; onboarding time halved; organisational productivity gains remain around 10% despite near-universal adoption. https://shiftmag.dev/this-cto-says-93-of-developers-use-ai-but-productivity-is-still-10-8013/

E

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), 1306–1308.https://doi.org/10.1126/science.adj0998

Erik Brynjolfsson, Danielle Li, Lindsey Raymond, Generative AI at Work, The Quarterly Journal of Economics, Volume 140, Issue 2, May 2025, Pages 889–942, https://doi.org/10.1093/qje/qjae044

European Central Bank. (2021). Hours worked in the euro area. ECB Economic Bulletin, Issue 6/2021, Article 1. European Central Bank.https://www.ecb.europa.eu/press/economic-bulletin/articles/2021/html/ecb.ebart202106_01~9c1a646a58.en.html

European Parliament. (2023, June 1). EU AI Act: First regulation on artificial intelligence. European Parliament Topics.https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence

European Parliament. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

EY. (2024, November). EY Work Reimagined Survey 2024: GenAI adoption at work. https://www.ey.com/en_cy/newsroom/2024/11/news-release-ey-reimagined-survey-2024

F

Fabrizio Dell'Acqua , Edward McFowland , Ethan Mollick , Hila Lifshitz , Katherine C. Kellogg , Saran Rajendran , Lisa Krayer , François Candelon , Karim R. Lakhani Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4921696

Fang, L., Yuan, Z., Zhang, K., Donati, D., & Sarvary, M. (2025). Generative AI and firm productivity: Field experiments in online retail (arXiv Working Paper No. 2510.12049).https://arxiv.org/abs/2510.12049

Federal Reserve Bank of New York. (2026). The labor market for recent college graduates. Federal Reserve Bank of New York.https://www.newyorkfed.org/research/college-labor-market

Federal Reserve Bank of New York. (n.d.). The labor market for recent college graduates. Federal Reserve Bank of New York.https://www.newyorkfed.org/research/college-labor-market#--:explore:unemployment

Federal Reserve Bank of St. Louis. (n.d.-a). Labor market [Data category]. FRED Economic Data.https://fred.stlouisfed.org/categories/11

Federal Reserve Bank of St. Louis. (n.d.-b). Employment level – Agriculture and related industries (Series LNS12034560) [Data series]. FRED Economic Data.https://fred.stlouisfed.org/data/LNS12034560

Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal, 42(12), 2195–2217. https://doi.org/10.1002/smj.3286

Florence Jaumotte, Jaden Kim, David Koll, Elmer Li, Longji Li, Giovanni Melina, Alina Song, and Marina Mendes Tavares. "Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age", Staff Discussion Notes 2026, 001 (2026), accessed 14/04/2026, https://doi.org/10.5089/9798229028196.006

Forrester. (2024). Your employees aren't ready for AI — prepare them with AIQ.https://www.forrester.com/blogs/your-employees-arent-ready-for-ai-prepare-them-with-aiq/

G

Gartner (March 2025). Agentic AI will autonomously resolve 80% of common customer service issues without human intervention by 2029; 30% reduction in operational costs.https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner-predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290

Gartner. (2024a). Gartner predicts conversational AI will reduce contact centre agent labour costs by $80 billion in 2026. Gartner Research.https://www.gartner.com/en/newsroom

Gartner. (2024b). Gartner predicts AI will resolve 80 percent of common customer service issues without human intervention by 2029. Gartner Newsroom.https://www.gartner.com/en/newsroom/press-releases/2024-03-25-gartner-predicts-ai-will-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-2029

Gassam, J. (2025, June 24). 92 million jobs gone: Who will AI erase first? Forbes. https://www.forbes.com/sites/janicegassam/2025/06/24/92-million-jobs-gone-who-will-ai-erase-first/

Georgieva, K. (2026, January 14). New skills and AI are reshaping the future of work. IMF Blog. https://www.imf.org/en/blogs/articles/2026/01/14/new-skills-and-ai-are-reshaping-the-future-of-work

Gibson, C. (2010). American demographic history chartbook: 1790 to 2000. demographicchartbook.com.https://demographicchartbook.com/wp-content/uploads/2015/11/Gibson-DemographicChartbook.pdf

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. Academy of Management Annals, 14(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Goh, E., Gallo, R., Hom, J., Strong, E., Weng, Y., Kerman, J., Cool, J. A., Kanjee, Z., Parsons, A. S., Ahuja, N., Horvitz, E., Yang, D., Milstein, A., Olson, A. P. J., Rodman, A., & Chen, J. H. (2024). Large language model influence on diagnostic reasoning: A randomized clinical trial. JAMA Network Open, 7(10), e2440969.https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2825399

Goldman Sachs Global Investment Research. (2023). The potentially large effects of artificial intelligence on economic growth. Goldman Sachs.https://www.gspublishing.com/content/research/en/reports/2023/03/27/d64e052b-0f6e-45d7-967b-d7be35fabd16.pdf

H

Handshake. (2025). State of the graduate: Class of 2025. Handshake.https://joinhandshake.com/network-trends/class-of-2025-graduation/

Hansen, M. T., & von Oetinger, B. (2001, March). Introducing T-shaped managers: Knowledge management's next generation. Harvard Business Review, 79(3), 106–116.https://hbr.org/2001/03/introducing-t-shaped-managers-knowledge-managements-next-generation

Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3706598.3713778

Hartmann, J., Exner, Y., & Domdey, S. (2025). The power of generative marketing: Can generative AI create superhuman visual marketing content? International Journal of Research in Marketing, 42(1), 13–31.https://doi.org/10.1016/j.ijresmar.2024.09.002

Harvard Business Review. (2025). How AI is redefining managerial roles. Harvard Business Review.https://hbr.org/2025/07/how-ai-is-redefining-managerial-roles

Hoffmann, M., Boysel, S., Nagle, F., Peng, S., & Xu, K. (2025). Generative AI and the nature of work (Harvard Business School Working Paper No. 25-021; CESifo Working Paper Series No. 11479; SSRN 5007084). Harvard Business School.https://www.hbs.edu/faculty/Pages/item.aspx?num=66593

Hu, K. (2023, February 2). ChatGPT sets record for fastest-growing user base. Reuters.https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/

Huang, J. (2026, January 21). Remarks at the World Economic Forum annual meeting [Panel discussion]. World Economic Forum, Davos, Switzerland.https://blogs.nvidia.com/blog/davos-wef-blackrock-ceo-larry-fink-jensen-huang/

Hui, X., Reshef, O., & Zhou, L. (2023). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market [Working paper].https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321

Humlum, A., & Vestergaard, E. (2025). Large language models, small labor market effects (NBER Working Paper No. 33777). National Bureau of Economic Research.https://doi.org/10.3386/w33777

I

IEA (2025), Energy and AI, IEA, Paris https://www.iea.org/reports/energy-and-ai, Licence: CC BY 4.0

IEA — International Energy Agency (April 2025). Energy and AI. Data centre electricity demand to double to 945 TWh by 2030; 85%+ of new capacity in US, China and EU.https://www.iea.org/reports/energy-and-ai

IGM Forum. (2023). AI and growth. University of Chicago Booth School of Business, Kent A. Clark Center for Global Markets.https://kentclarkcenter.org/surveys/ai-and-growth/

IMF — International Monetary Fund (January 2024). Gen-AI: Artificial Intelligence and the Future of Work. Staff Discussion Note SDN2024/001. 40% of global jobs exposed; 60% in advanced economies.https://www.imf.org/en/blogs/articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity

Internal Revenue Service. (2026, February 19). Federal unemployment tax. https://www.irs.gov/individuals/international-taxpayers/federal-unemployment-tax

Internal Revenue Service. (2026, January 20). Topic no. 751: Social Security and Medicare withholding rates. https://www.irs.gov/taxtopics/tc751

International Energy Agency. (2025a). Energy and AI. IEA.https://www.iea.org/reports/energy-and-ai

International Energy Agency. (2025b). Electricity 2025: Analysis and forecast to 2027. IEA Publications.https://www.iea.org/reports/electricity-2025

International Federation of Robotics. (2025, September 25). World Robotics 2025 report – industrial robots – released by IFR: Global robot demand in factories doubles over 10 years.https://ifr.org/ifr-press-releases/news/global-robot-demand-in-factories-doubles-over-10-years

International Labour Organization. (2023). The future of work in the textile, clothing, leather and footwear sectors. ILO Sectoral Policies Department.https://www.ilo.org/publications/future-work-textiles-clothing-leather-and-footwear

International Monetary Fund. (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001. https://www.imf.org/-/media/files/publications/sdn/2024/english/sdnea2024001.pdf

International Telecommunication Union. (2024). Measuring digital development: Facts and figures 2024. ITU Publications.https://www.itu.int/en/mediacentre/Pages/PR-2024-11-27-facts-and-figures.aspx

Introducing T-Shaped Managers: Knowledge Management's Next Generation by Morten T. Hansen and Bolko von Oetinger https://hbr.org/2001/03/introducing-t-shaped-managers-knowledge-managements-next-generation

---

J

Jaumotte and others. 2026. Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age. Staff Discussion Note SDN2026/001.International Monetary Fund, Washington D.C. https://www.imf.org/-/media/files/publications/sdn/2026/english/sdnea2026001.pdf

Jaumotte, F., Kim, J., Koll, D., Li, E., Li, L., Melina, G., Song, A., & Mendes Tavares, M. (2026). Bridging skill gaps for the future: New jobs creation in the AI age (Staff Discussion Notes No. 2026/001). International Monetary Fund.https://doi.org/10.5089/9798229028196.006

Jose, B., Joseph, D., Mohan, V., Alexander, E., Varghese, S. K., & Roy, A. (2025). Outsourcing cognition: the psychological costs of AI-era convenience. Frontiers in psychology, 16, 1645237. https://doi.org/10.3389/fpsyg.2025.1645237

K

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. Labour Economics, 17(2), 303–316.https://doi.org/10.1016/j.labeco.2009.09.002

Katelynn Harris, "Forty years of falling manufacturing employment ," Beyond the Numbers: Employment & Unemployment, vol. 9 / no. 16 (U.S. Bureau of Labor Statistics, November 2020), https://www.bls.gov/opub/btn/volume-9/forty-years-of-falling-manufacturing-employment.html

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.https://doi.org/10.2307/2118323

Klarna / CX Dive (May 2025). Klarna pivots back to human agents after overpivoting on AI cost-cutting.https://www.customerexperiencedive.com/news/klarna-reinvests-human-talent-customer-service-AI-chatbot/747586/

Klarna / Yahoo Finance (November 2025). AI agent now equivalent to 853 employees; response times improved 82%; repeat issues down 25%.https://finance.yahoo.com/news/klarna-says-ai-agent-doing-070000869.html

Klarna. (2024, February 27). Klarna AI assistant handles two-thirds of customer service chats in its first month. Klarna Newsroom.https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task. arXiv preprint arXiv:2506.08872.https://arxiv.org/abs/2506.08872

Kumar, P. S. (2024). Technostress: A comprehensive literature review on dimensions, impacts, and management strategies. Computers in Human Behavior Reports, 16, Article 100475.https://doi.org/10.1016/j.chbr.2024.100475

L

Lebergott, S. (1966). Labor force and employment, 1800–1960. In D. S. Brady (Ed.), Output, employment, and productivity in the United States after 1800 (pp. 117–204). National Bureau of Economic Research.https://www.nber.org/system/files/chapters/c1567/c1567.pdf

Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2021).Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What1. MIS Quarterly 1 September 2021; 45 (3): 1501–1526. https://doi.org/10.25300/MISQ/2021/16564

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. Organization Science, 33(1), 126–148. https://doi.org/10.1287/orsc.2021.1549

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence

effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25). ACM. https://doi.org/10.1145/3706598.3713778

Lee, J., Todri, V., Adamopoulos, P., & Ghose, A. (2025). Generative AI and creative content: A field study of human and AI-created advertisements. SSRN.https://doi.org/10.2139/ssrn.5638311

Levanon, G., Sigelman, M., Mamertino, M., de Zeeuw, M., & Guilford, G. (2025, July). No country for young grads: The structural forces that are reshaping entry-level work. Burning Glass Institute.https://www.burningglassinstitute.org/research/no-country-for-young-grads

Levy Yeyati, E. (2025, July 10). Hybrid jobs: How AI is rewriting work in finance. Brookings Institution.https://www.brookings.edu/articles/hybrid-jobs-how-ai-is-rewriting-work-in-finance/

Li, C., Wang, H., Jiang, S., & Gu, B. (2024). The effect of AI-enabled credit scoring on financial inclusion: Evidence from an underserved population of over one million. MIS Quarterly, 48(4), 1803–1834.https://doi.org/10.25300/MISQ/2024/18340

Lisa B. Kahn, The long-term labor market consequences of graduating from college in a bad economyhttps://doi.org/10.1016/j.labeco.2009.09.002

Liu, T. L., Hetherington, T. C., Dharod, A., Carroll, T., Bundy, R., Nguyen, H., Bundy, H. E., Isreal, M., McWilliams, A., & Cleveland, J. A. (2024). Does AI-powered clinical documentation enhance clinician efficiency? A longitudinal study. NEJM AI, 1(12).https://doi.org/10.1056/AIoa2400659

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1–16). Association for Computing Machinery. https://doi.org/10.1145/3313831.3376727

Lyel l, D., & Coiera, E. (2017). Automation bias and verification complexity: A systematic review. Journal of the American Medical Informatics Association, 24(2), 423–431.https://doi.org/10.1093/jamia/ocw105

M

Mallon, D., Cantrell, S., Forsythe, J., et al. (2025, March). Closing the experience gap through talent development. In 2025 Global Human Capital Trends. Deloitte Insights.https://www.deloitte.com/us/en/insights/topics/talent/human-capital-trends/2025/closing-the-experience-gap-through-talent-development.html

Manuel Hoffmann, Sam Boysel, Frank Nagle, Sida Peng and Kevin Xu (2025). Generative AI and the Nature of Work. https://www.hbs.edu/faculty/Pages/item.aspx?num=66593

Massenkoff, M., & McCrory, P. (2026, March 5). Labor market impacts of AI: A new measure and early evidence. Anthropic.
https://cdn.sanity.io/files/4zrzovbb/website/2b5bbaf2c1eb81dbf6e6fb813c1a24e35a64d376.pdf

Mayer, H., Yee, L., Chui, M., & Roberts, R. (2025, January 28). Superagency in the workplace: Empowering people to unlock AI's full potential. McKinsey & Company.https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work

McCoy, L. G., Swamy, R., Sagar, N., Wang, M., Bacchi, S., Fong, J. M. N., Tan, N. C. K., Tan, K., Buckley, T. A., Brodeur, P., Celi, L. A., Manrai, A. K., Humbert, A., & Rodman, A. (2025). Assessment of large language models in clinical reasoning: A novel benchmarking study. NEJM AI, 2(10).https://doi.org/10.1056/AIdbp2500120

McKinsey Global Institute. (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. McKinsey & Company.https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages

McKinsey Global Institute. (2023a). The economic potential of generative AI: The next productivity frontier. McKinsey & Company.https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier

McKinsey Global Institute. (2023b). Generative AI and the future of work in America. McKinsey & Company.https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america

McKinsey Global Institute. (2025). The state of AI in 2025. McKinsey & Company.https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

McKinsey Health Institute & World Economic Forum (January 2026). The Human Advantage: Stronger Brains in the Age of AI. Brain capital defined as the combined value of brain health and brain skills; scaling existing interventions could unlock $6.2 trillion in GDP gains by 2050; investment in brain capital framed as a strategic economic imperative in the AI era. https://www.mckinsey.com/mhi/our-insights/the-human-advantage-stronger-brains-in-the-age-of-ai

Microsoft. (2024, August 8). DAX Copilot: New customization options and AI capabilities for even greater productivity. Microsoft Industry Blog.https://www.microsoft.com/en-us/industry/blog/healthcare/2024/08/08/dax-copilot-new-customization-options-and-ai-capabilities-for-even-greater-productivity/

Mintzberg, H. (1980). Structure in 5's: A synthesis of the research on organization design. Management Science, 26(3), 322–341. https://doi.org/10.1287/mnsc.26.3.322

Muehlemann, Samuel & Wolter, Stefan C., 2011. "Firm-sponsored training and poaching externalities in regional labor markets," Regional Science and Urban Economicshttps://ideas.repec.org/a/eee/regeco/v41y2011i6p560-570.html

---

N

National Association of Colleges and Employers. (2023). The class of 2023: Inequity continues to underpin internship participation and pay status. NACE.https://www.naceweb.org/diversity-equity-and-inclusion/trends-and-predictions/the-class-of-2023-inequity-continues-to-underpin-internship-participation-and-pay-status/

National Association of Colleges and Employers. (2025, August 26). Average starting salary for Class of 2024 shows mild gain.https://www.naceweb.org/job-market/compensation/average-starting-salary-for-class-of-2024-shows-mild-gain

National Bureau Of Economic Research, Applying Ai To Rebuild Middle Class Jobs, David Autor, Http://Www.Nber.Org/Papers/W32140

National Bureau Of Economic Research, Still Waters, Rapid Currents: Early Labor Market Transformation Under Generative Ai, Anders Humlum, Emilie Vestergaard, http://Www.Nber.Org/Papers/W33777

National Bureau Of Economic Research, The Simple Macroeconomics Of AI, Daron Acemoglu https://ideas.repec.org/a/eee/regeco/v41y2011i6p560-570.html

National Center for Education Statistics. (n.d.). Annual earnings by educational attainment. Condition of Education. U.S. Department of Education.https://nces.ed.gov/programs/coe/indicator/cba/annual-earnings

National Institute of Standards and Technology. (2024). AI Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce.https://doi.org/10.6028/NIST.AI.100-1

NBAA – National Business Aviation Association. (2025, November/December). When tech fails, analog gear and skills can save you. Business Aviation Insider.https://nbaa.org/news/business-aviation-insider/2025/when-tech-fails-analog-gear-and-skills-can-save-you/

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

O

OECD. (2024). OECD employment outlook 2024: The net zero transition and the labour market. OECD Publishing.https://www.oecd.org/en/publications/oecd-employment-outlook-2024_ac97c260-en.html

OECD. (2025). OECD skills outlook 2025: Building the skills of the 21st century for all. OECD Publishing. https://doi.org/10.1787/26163cd3-en

OECD. (2026, January). AI use by individuals surges across the OECD as adoption by firms continues to expand. OECD.https://www.oecd.org/en/about/news/announcements/2026/01/ai-use-by-individuals-surges-across-the-oecd-as-adoption-by-firms-continues-to-expand.html

Ong, A. Y., Merle, D. A., Pollreisz, A., Wagner, S. K., Sevgi, M., Keane, P. A., Huemer, R., Oehling, J., Jäger, M., & Huemer, J. (2026). Flight rules for clinical AI: Lessons from aviation for human-AI collaboration in medicine. npj Digital Medicine, 9, 201.https://doi.org/10.1038/s41746-026-02410-1

Oreopoulos, P., von Wachter, T., & Heisz, A. (2012). The short- and long-term career effects of graduating in a recession. American Economic Journal: Applied Economics, 4(1), 1–29.https://pubs.aeaweb.org/doi/pdfplus/10.1257/app.4.1.1

P

Pachidi, S., Berends, H., Faraj, S., & Huysman, M. (2021). Make way for the algorithms: Symbolic actions and change in a regime of knowing. Organization Science, 32(1), 18–41. https://doi.org/10.1287/orsc.2020.1377

Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv preprint arXiv:2302.06590.https://arxiv.org/pdf/2302.06590

Perna, M. C. (2025, March 24). New data: 41% of Gen Z workers are sabotaging their employers' AI strategy. Forbes.https://www.forbes.com/sites/markcperna/2025/03/24/new-data-41-of-gen-z-workers-are-sabotaging-their-employers-ai-strategy/

Polanyi, M. (2009). The tacit dimension (A. Sen, Foreword). University of Chicago Press. (Original work published 1966)https://press.uchicago.edu/ucp/books/book/chicago/T/bo6035368.html

Power of generative marketing: Can generative AI create superhuman visual marketing content? Jochen Hartman, Yannick Exner, Samuel Domdey https://doi.org/10.1016/j.ijresmar.2024.09.002

PricewaterhouseCoopers. (2024). Global Workforce Hopes and Fears Survey 2024. PwC.https://www.pwc.com/gx/en/issues/workforce/hopes-and-fears-2024.html

PwC. (2025, June 3). AI linked to a fourfold increase in productivity growth and 56% wage premium [Press release].https://www.pwc.com/gx/en/news-room/press-releases/2025/ai-linked-to-a-fourfold-increase-in-productivity-growth.html

R

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. Academy of Management Review, 46(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Ramnath, S., Tasci, M., & Zabek, M. (2025, November 24). Are young college graduates losing their edge in the job market? Economic Commentary (2025-14). Federal Reserve Bank of Cleveland.https://www.clevelandfed.org/publications/economic-commentary/2025/ec-202514-are-young-college-graduates-losing-their-edge-in-the-job-market

Randazzo, Steven and Lifshitz-Assaf, Hila and Kellogg, Katherine and Dell'Acqua, Fabrizio and Mollick, Ethan R. and Candelon, François and Lakhani, Karim R., Cyborgs, Centaurs and Self-Automators: The Three Modes of Human-GenAI Knowledge Work and Their Implications for Skilling and the Future of Expertise (December 08, 2025). The Wharton School Research Paper, Harvard Business School Working Paper No. 26-036, Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 26-036, http://dx.doi.org/10.2139/ssrn.4921696

Randstad. (2025, September 22). 24% decline in financial services junior roles spurs shift to an AI-driven skills-first revolution.https://www.randstad.com/press/2025/24-decline-financial-services-junior-roles-spurs-shift-to-an-ai-driven-skills-first/

Revelio Labs (December 2025). 2025 Workforce Insights Wrapped. Highly AI-exposed entry-level roles declined over 40% since January 2023.https://www.reveliolabs.com/news/social/2025-workforce-insights-wrapped/

Revelio Labs (July 2025). Entry-level job postings 50% below January 2022 levels.https://www.reveliolabs.com/news/macro/jobs-outlook-july-2025/

Revelio Labs / Bloomberg (August 2025). AI exposure associated with 11% drop in entry-level demand and 7% increase for senior roles in same occupations.https://www.reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-unemployment/

Revelio Labs. (2024). The AI job market: Trends in hiring and displacement. Revelio Labs Workforce Intelligence Report.https://www.reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-unemployment/

Rice, L. (2025, July 29). ARM Institute publishes new future skill requirements report. ARM Institute. https://arminstitute.org/news/future-skill-report/

Risk Management Association. (2024). Modeling the future: RMA survey of model risk management.https://www.rmahq.org/journal-articles/2024/october-november-2024/in-rma-s-model-risk-management-survey-a-picture-of-banks-diligence-and-frustrations/

S

Selbst, A. D., boyd, d., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. Proceedings of the Conference on Fairness, Accountability, and Transparency, 59–68.https://dl.acm.org/doi/10.1145/3287560.3287598

Shakked Noy, Whitney Zhang ,Experimental evidence on the productivity effects of generative artificial intelligence.Science381,187-192(2023). DOI:10.1126/science.adh2586

Shen, J. H., & Tamkin, A. (2026). How AI impacts skill formation. https://doi.org/10.48550/arXiv.2601.20245

Shrestha, Y. R., Ben-Menahem, S. M., & von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. California Management Review, 61(4), 66–83. https://doi.org/10.1177/0008125619862257

Singla, A., Sukharevsky, A., & Yee, L. (2025). The state of AI: How organizations are rewiring to capture value. McKinsey & Company.https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Sivakumar, R., Lue, B., & Kundu, S. (2025). FDA approval of artificial intelligence and machine learning devices in radiology: A systematic review. JAMA Network Open, 8(11), e2542338.https://pmc.ncbi.nlm.nih.gov/articles/PMC12595527/

Snow, C. C., & Ketchen, D. J. (2014). Typology-driven theorizing: A response to Delbridge and Fiss. Academy of Management Review, 39(2), 231–233. https://doi.org/10.5465/amr.2013.0388

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. Science, 333(6043), 776–778.https://doi.org/10.1126/science.1207745

Spencer Stuart. (2026, February). 2025 S&P 1500 CEO transitions: Behind the CEO moment. https://www.spencerstuart.com/research-and-insight/2025-sp-1500-ceo-transitions-behind-the-ceo-moment

Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. Computers in Human Behavior, 160, Article 108386.https://doi.org/10.1016/j.chb.2024.108386

Stanford Digital Economy Lab. (2025). AI and the labor market: Early evidence from payroll data. Stanford University.https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf

Stanford HAI (April 2025). AI Index Report 2025. US private AI investment $109.1 billion in 2024 — nearly 12× China's $9.3 billion.https://hai.stanford.edu/ai-index/2025-ai-index-report

Stanford HAI (April 2026). AI Index Report 2026. US private AI investment $285.9 billion in 2025 — 23× China. US hosts 5,427 data centres.https://hai.stanford.edu/ai-index/2026-ai-index-report

Stanford University Human-Centered AI Institute. (2024). AI Index Report 2024. Stanford University.https://aiindex.stanford.edu/report/

Stevenson, C., Theissing, H., & de Wit, J. (2022). LLMs and creativity: Evaluating AI writing on the Alternative Uses Task [Working paper].https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4285244

Sukharevsky, A., Kerr, D., Hjartar, K., Hämäläinen, L., Bout, S., & Di Leo, V. (2025, June 13). Seizing the agentic AI advantage. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage

Sukharevsky, A., Krivkovich, A., Gast, A., Storozhev, A., Maor, D., Mahadevan, D., Hämäläinen, L., & Durth, S. (2025, September 26). The agentic organization: Contours of the next paradigm for the AI era. McKinsey & Company.https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era

---

T

The Impact of Visual Generative AI on Advertising Effectiveness, Hyesoo Lee, Vilma Todri, Panagiotis Adamopoulos, Anindya Ghose https://doi.org/10.2139/ssrn.5638311

The paradox of rising investment and elusive returns, Richard Horton, Jan Michalski, Stacey Winters, Douglas Gunn (2025) https://www.deloitte.com/global/en/issues/generative-ai/ai-roi-the-paradox-of-rising-investment-and-elusive-returns.html

Thompson, J. D. (1967). Organizations in action: Social science bases of administrative theory. McGraw-Hill.

Tyna Eloundou et al. ,GPTs are GPTs: Labor market impact potential of LLMs.Science384,1306-1308(2024).DOI:10.1126/science.adj0998

U

U.S. Bureau of Labor Statistics, Current Employee Statistics https://www.bls.gov/ces/

U.S. Bureau of Labor Statistics, Forty years of falling manufacturing employment https://www.bls.gov/opub/btn/volume-9/forty-years-of-falling-manufacturing-employment.html

U.S. Bureau of Labor Statistics, Market Research Analysts https://www.bls.gov/ooh/business-and-financial/market-research-analysts.htm

U.S. Bureau of Labor Statistics. (2025). Market research analysts: Occupational outlook handbook. U.S. Department of Labor.https://www.bls.gov/ooh/business-and-financial/market-research-analysts.htm

U.S. Department of Labor, Employment and Training Administration. (n.d.). WARN Act compliance assistance. https://www.dol.gov/agencies/eta/layoffs/warn

U.S. Department of Labor. (n.d.). ONET OnLine: 13-1161.01 — Search marketing strategists*.https://www.onetonline.org/link/summary/13-1161.01

U.S. Department of Labor. (n.d.). ONET OnLine: 13-2041.00 — Credit analysts*.https://www.onetonline.org/link/summary/13-2041.00

U.S. Department of Labor. (n.d.). ONET OnLine: 13-2072.00 — Loan officers*.https://www.onetonline.org/link/summary/13-2072.00

U.S. Department of Labor. (n.d.). ONET OnLine: 29-1216.00 — General internal medicine physicians*.https://www.onetonline.org/link/summary/29-1216.00

U.S. Department of Labor. (n.d.). Plant closings and layoffs (WARN Act). U.S. Department of Labor.https://www.dol.gov/general/topic/termination/plantclosings

U.S. Food and Drug Administration. (2025). Artificial intelligence and machine learning (AI/ML)-enabled medical devices.https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices

United Nations, Department of Economic and Social Affairs, Population Division. (2024). World population prospects 2024 [Data set]. United Nations.https://population.un.org/wpp/downloads?folder=Standard%20Projections&group=Population

Upstart Holdings, Inc. (2025). Annual report (Form 10-K, period ending December 31, 2024). U.S. Securities and Exchange Commission.https://www.sec.gov/Archives/edgar/data/1647639/000164763925000018/upst-20241231.htm

---

V

Vaccaro, M., Almaatouq, A. & Malone, T. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nat Hum Behav 8, 2293–2303 (2024). https://doi.org/10.1038/s41562-024-02024-1

Vaccaro, M., Almaatouq, A., & Malone, T. W. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour, 8(12), 2293–2303.https://doi.org/10.1038/s41562-024-02024-1

Van den Broek, E., Sergeeva, A., & Huysman, M. (2021).When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring1. MIS Quarterly 1 September 2021; 45 (3): 1557–1580. https://doi.org/10.25300/MISQ/2021/16559

Vietnam Textile and Apparel Association (VITAS), cited in IMARC Group (2025). Around 20% of Vietnamese textile companies had implemented Industry 4.0 technologies by end of 2023, with an expected rise to 40% by 2025; rising wages accelerating the shift to automation.https://www.imarcgroup.com/vietnam-textile-market

Voth, H.-J. (2001). The longest years: New estimates of labor input in England, 1760–1830. Journal of Economic History, 61(4), 1065–1082.https://www.crei.cat/wp-content/uploads/users/working-papers/voth_longestyears.pdf

---

W

Webb, M. (2020). The impact of artificial intelligence on the labor market [Working paper]. Stanford University. https://www.michaelwebb.co/

Williamson, O. E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. Administrative Science Quarterly, 36(2), 269–296. https://doi.org/10.2307/2393356

Wilson, H. J., & Daugherty, P. R. (2024, September). Embracing gen AI at work. Harvard Business Review. https://hbr.org/2024/09/embracing-gen-ai-at-work

World Economic Forum. (2020). The future of jobs report 2020. World Economic Forum. https://www3.weforum.org/docs/WEF_Future_of_Jobs_2020.pdf

World Economic Forum. (2023). The future of jobs report 2023. World Economic Forum. https://www3.weforum.org/docs/WEF_Future_of_Jobs_2023.pdf

World Economic Forum. (2025). The future of jobs report 2025. World Economic Forum.https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf

World Economic Forum. (January 2025). Future of Jobs Report 2025. 59% of the global workforce will require reskilling or upskilling by 2030; 39% of core skills will change; analytical thinking, creative thinking, resilience and adaptability among the most critical capabilities.https://www.weforum.org/press/2025/01/future-of-jobs-report-2025-78-million-new-job-opportunities-by-2030-but-urgent-upskilling-needed-to-prepare-workforces/

World Economic Forum. (January 2025). Future of Jobs Report 2025. 92 million jobs displaced by 2030; 170 million new roles created; net gain of 78 million. AI affects 86% of businesses.https://www.weforum.org/press/2025/01/future-of-jobs-report-2025-78-million-new-job-opportunities-by-2030-but-urgent-upskilling-needed-to-prepare-workforces/

Y

Yin, M., Vu, H., & Persico, C. (2026). How (un)stable are LLM occupational exposure scores? Evidence from multi-model replication (NBER Working Paper No. 35110). National Bureau of Economic Research. https://doi.org/10.3386/w35110

Z

Zest AI. (2024). How AI is revolutionizing credit union lending amid rising delinquencies. Zest AI Blog.https://www.zest.ai/learn/blog/much-ado-about-delinquencies/

Zest AI. (2025, November 4). Zest AI secures strategic investment from customers in oversubscribed round. Zest AI Press Release.https://www.zest.ai/company/announcements/zest-ai-secures-strategic-investment-from-customers-in-oversubscribed-round

APPENDIX

Appendix A: Methodological Breakdown of the Post-AI Pay Structure

This appendix formalizes the model used to construct Figure 2. The question it is designed to answer is narrow but important: if AI compresses the entry-level and early-career layers of knowledge work, do the resulting payroll savings necessarily translate into durable labor-cost reduction for firms? The conclusion from the final calibrated model is that they do not necessarily do so. Under the parameter values used here, savings from a thinner junior base are largely or entirely offset once firms still require a stable experienced workforce and must pay a premium for increasingly scarce mid-level, senior, and executive talent.

The appendix should be read as an illustrative exercise designed to prove that assumptions companies may take as given are not necessarily true, rather than a forecast. Every parameter approximates what the data suggests for illustrative value, not a claim about the uniquely correct value for every firm. In several places the rounded figure shown in the model is best understood as an approximate midpoint of a reasonable range. The goal is directional realism rather than absolute accuracy.

Step 1. Fix the unit of analysis and the five workforce layers

The unit of analysis is a representative 100-person white-collar, knowledge-work firm. The model uses five workforce layers: entry-level, early-career, mid-level, senior, and executive. Public labor datasets do not directly report this exact five-rung hierarchy, so the model combines direct labor-market evidence with transparent calibration choices.

Step 2. Calibrate the baseline headcount split

The baseline structure is set at 18 / 22 / 43 / 13 / 4. This is an approximate measure of a workforce split as there is no clear existing hierarchy that maps across all relevant firms in the sector. The table below shows how each layer was chosen and what evidence most directly supports it.

Level	Headcount per 100 workers	Rationale for use	Source / rough benchmark
Entry-level	18	Modeled as a relatively narrow first rung. Knowledge occupations have sufficiently long tenure that despite the high volume within the entry-level, this is a more modest share. This balances between many knowledge occupations with smaller entry-level teams, and some in professional services with a significant entry-level force.	BLS tenure data: median tenure is 4.8 years in management, professional, and related occupations, 4.5 years in business and financial operations, and 4.3 years in computer and mathematical occupations (Bureau of Labor Statistics, 2024).

Early-career	22	Anchored as part of a combined junior pipeline of 40 workers per 100. With entry-level fixed at 18 as a relatively narrow first rung, early-career is set at 22 to create a meaningful second pipeline layer without making the model resemble an exaggerated professional-services pyramid. This keeps the first two layers substantial but not dominant in a firm where a broad experienced middle still carries most output.	BLS tenure medians in knowledge occupations — 4.8 years in management, professional, and related occupations, 4.5 years in business and financial operations, and 4.3 years in computer and mathematical occupations — support a sizable but not majority early-stage workforce (Bureau of Labor Statistics, 2024).
Mid-level	43	Set as the residual broad middle once the combined junior pipeline and upper layers are fixed. With 40 workers in entry-level plus early-career roles and 17 in senior plus executive roles, the remaining 43 workers form the dominant experienced, non-senior core of the organization. This is consistent with a knowledge-work hierarchy in which leadership is thin and most labor sits below senior leadership but above the junior pipeline.	BLS OEWS shows Management Occupations at 12.18 percent in Finance and Insurance and 12.59 percent in Professional, Scientific, and Technical Services, while Top Executives are 3.88 percent and 4.39 percent, implying a thin upper hierarchy and leaving the broad middle as the dominant layer (Bureau of Labor Statistics, 2023b; 2023c).
Senior	13	Semi-anchored to observed formal management density, then widened modestly for senior individual contributors and director-like roles that BLS does not separate cleanly.	Management Occupations are 12.18% in Finance and Insurance and 12.59% in Professional, Scientific, and Technical Services; Top Executives are 3.88% and 4.39%, implying roughly 8 to 9 non-top managers per 100 workers before adding senior ICs (Bureau of Labor Statistics, 2023b; 2023c).

Executive	4	Based on clear statistics on share of workforce in top executive roles	BLS OEWS reports Top Executives at 3.64% of employment in Information, 3.88% in Finance and Insurance, and 4.39% in Professional, Scientific, and Technical Services. A value of 4 per 100 workers is the midpoint of that observed range (Bureau of Labor Statistics, 2023a; 2023b; 2023c).

Taken together, these values imply a baseline with a thin executive tier, a modest senior layer, and a broad middle beneath leadership. That is a more defensible starting point for a paper concerned with the broad white-collar economy than either a flat corporate hierarchy or an exaggerated professional-services pyramid.

Step 3. Set wage anchors for each layer

Each wage input is chosen to approximate average pay figures for knowledge workers across sectors. Where public data provide a direct benchmark, that benchmark is used as the anchor. Where no public series maps cleanly onto the rung, the value shown is an approximation anchored between adjacent benchmarks and the broader occupational evidence.

Level	Wage used	Underlying benchmark	Rationale for use	Source
Entry-level	$65,000	NACE reports an average starting salary of $65,677 for the Class of 2024.	Best fit for true entry-level professional pay because it measures recent graduate starting salaries rather than a broad age band that includes workers well past entry level.	National Association of Colleges and Employers, 2025

Early-career	$80,000	Two useful reference points bracket this layer: NCES reports median annual earnings of $66,600 for 25- to 34-year-olds with a bachelor's degree working full time, year round in 2022, while BLS reports a median annual wage of $80,920 for business and financial occupations in May 2024.	No public series isolates the second rung of a white-collar hierarchy. The model therefore uses $80,000 as an approximation between a lower young-bachelor earnings anchor and a broader white-collar occupational median, keeping early-career pay above the true starting-salary benchmark and below the mid-level manager proxy.	National Center for Education Statistics, n.d.; Bureau of Labor Statistics, 2025a
Mid-level	$105,000	BLS reports a median annual wage of $102,950 for general and operations managers in May 2024.	Approximates middle-management and experienced team-lead pay more credibly than the broader all-management median.	Bureau of Labor Statistics, 2025b
Senior	$135,000	In Professional, Scientific, and Technical Services, general and operations managers had a median hourly wage of $64.85, which annualizes to about $134,888 at 2,080 hours.	Approximates a sector-specific senior-manager benchmark for knowledge services rather than an economy-wide average.	Bureau of Labor Statistics, 2023c

Executive	$215,000	BLS reports top Executives in Management of Companies and Enterprises had an annual mean wage of $217,070.	Approximates the top rung for executives in knowledge-work firms.	Bureau of Labor Statistics, 2025b; 2023d

Step 4. Specify the AI-compressed headcount structure

The compressed case changes only the bottom two layers. Entry-level headcount falls from 18 to 12 and early-career headcount falls from 22 to 17, while mid-level, senior, and executive headcount remain constant. This follows the core idea that early career stages are most impacted by AI and that any task-level replacement of more senior roles would be at least partly offset by increased productivity from AI tools.

Level	Baseline headcount	AI-compressed headcount	Rationale for change or non-change	Source
Entry-level	18	10	Represents substantial compression at the true entry point of the pipeline. This is directionally aligned with evidence of deteriorating entry-level employment outcomes in AI-exposed work.	Brynjolfsson et al., 2025; Federal Reserve Bank of New York, 2026
Early-career	22	16	Represents a somewhat smaller but still material contraction in the second pipeline layer, consistent with weaker demand for junior professional labor.	Brynjolfsson et al., 2025; Federal Reserve Bank of New York, 2026

Mid-level	43	43	Held constant because firms still require experienced workers to supervise, interpret, and execute more complex workflows after AI absorbs lower-level tasks.	Dell'Acqua et al., 2023; Davis, 2026
Senior	13	13	Held constant because senior judgment, accountability, and tacit knowledge remain necessary even under AI augmentation.	Davis, 2026; Acemoglu & Restrepo, 2019
Executive	4	4	Held constant because executive demand is driven less by routine task volume and more by organizational oversight, coordination, and strategic direction.	Acemoglu & Restrepo, 2019; Davis, 2026

Step 5. Apply experienced-worker scarcity and complementarity premia

The model then applies scarcity and complementarity premia of 15 percent for mid-level workers, 10 percent for senior workers, and 5 percent for executives. These are scenario parameters, not direct empirical estimates. They are justified by three literatures: evidence that external hires can cost materially more than internal promotions; labor-economics evidence that tighter relative supply and stronger relative demand for skill raise wage differentials; and task-based automation literature showing that technology can displace some workers while increasing the value of complementary expertise.

Level	Premium used	Rationale for use	Source

Mid-level	20%	Set close to, Bidwell's observed 18 percent external-hire premium. This is the layer most plausibly exposed to a real market premium if firms underbuild junior pipelines and later need to buy experienced talent externally.	Bidwell, 2011; Katz & Murphy, 1992; Davis, 2026
Senior	15%	Set below the mid-level premium because not every senior role is filled through open-market hiring, but still high enough to reflect tighter supply and stronger returns to tacit knowledge.	Katz & Murphy, 1992; Acemoglu & Restrepo, 2019; Davis, 2026
Executive	10%	Set as the smallest premium because executive compensation is more idiosyncratic and shaped by internal succession, governance, and incentives. The point is to capture some top-rung tightening without overstating a weakly identified parameter.	Acemoglu & Restrepo, 2019; Davis, 2026

Step 6. Compute payroll and interpret the result

Total payroll is computed as the sum, across all five layers of headcount multiplied by wage. The table below shows headcount, average wage, and resulting payroll by layer in the baseline and compressed-premium cases.

Level	Baseline headcount	Baseline wage ($000s)	Baseline payroll	AI-compressed headcount	AI-compressed wage ($000s)	AI-compressed payroll
Entry-level	18	65.0	$1,170,000	10	65.0	$650,000
Early-career	22	80.0	$1,760,000	16	80.0	$1,280,000
Mid-level	43	105.0	$4,515,000	43	126	$5,418,000
Senior	13	135.0	$1,755,000	13	155.25	$2,018,250
Executive	4	215.0	$860,000	4	236.5	$946,000

Total	100		$10,060,000	86		$10,312,250

The result is the central conclusion of the model. Total headcount falls from 100 to 86, yet total payroll rises slightly from $10.06 million to $10.3 million. Average payroll per worker rises from $100,600 to $119,910. The model therefore supports a narrower and more defensible claim than simple layoff arithmetic would suggest: once firms still need to preserve output and the market for experienced talent tightens, AI-driven reductions in junior hiring may cease to function as a durable net cost reducer.

Step 7. State the conclusion and the caveats clearly

Three caveats are important. First, the baseline split is calibrated from multiple datasets and is not directly observed as a five-rung hierarchy in any single public source. Confidence is therefore highest at the executive layer, reasonably strong at the senior layer, and materially weaker in the exact boundary between early-career and mid-level. Second, the compression pattern and the 15 / 10 / 5 premium schedule are scenario assumptions, albeit informed ones, not causal estimates. Third, the model is designed for broad white-collar and knowledge-work settings. It should not be applied mechanically to sectors with very different labor structures, such as frontline services, healthcare delivery, or manufacturing. The appendix is not attempting to predict the exact future composition of the labor market. It is demonstrating, in a transparent and reproducible way, how a thinner junior pipeline can coexist with flat or slightly higher total payroll once experienced labor becomes scarcer and more valuable.

Appendix B: Methodological Walkthrough — Task Scoring and Sensitivity Analysis

The worked example in section 2.2.3 develops the SEO Marketing Strategist case in non-technical terms. This appendix preserves the full methodology for readers who need to verify, replicate, or extend the analysis. It covers the Eloundou exposure rubric, the published scores for marketing occupations, the mapping from exposure category to governance tier, and the calculations behind each productivity scenario.

B.1 The Eloundou exposure rubric

Eloundou et al. (2024), in their Science paper “GPTs are GPTs,” developed the task-level exposure rubric used in the worked example. Each task in the O*NET inventory receives one of three labels: E0 (AI cannot meaningfully help with this task), E1 (an LLM can do this task on its own), or E2 (an LLM can do this task when embedded in specialized software, with the combined system cutting task time by at least 50% at comparable quality).

Three aggregate measures follow from these labels. Direct exposure (E1) counts only E1 tasks at full weight: the floor representing what AI can do today without additional tooling. Intermediate exposure (E1 + 0.5 × E2) counts E1 tasks at full weight and E2 tasks at half weight: the realistic middle estimate reflecting the world most enterprises are in, where complementary tools exist but are not yet fully integrated. Full exposure (E1 + E2) counts both E1 and E2 at full weight: the upper bound representing maximum exposure when every available tool is in place. Applied to the entire U.S. economy, roughly 80% of the workforce holds occupations where at least 10% of tasks are affected under the intermediate measure; in a reversal of the historical pattern, higher-income knowledge-work occupations face greater AI exposure than lower-income ones.

B.2 Published scores for search marketing and related occupations

For Search Marketing Strategists (SOC 13-1161.01), applying the Eloundou rubric to the occupation's 36 O*NET tasks — with the full task-level scoring set out in 2.6 — yields:

- Direct exposure = 27.8% — 10 of 36 tasks rated E1 (LLM alone).
- Intermediate exposure = 62.5% — more than double Eloundou's all-occupation mean of 30%, placing the function firmly in the high-exposure range.
- Full exposure = 97.2% — 35 of the 36 listed tasks are potentially exposed once LLMs are deployed within complementary marketing software; the remaining task is rated E0.

(Related marketing occupations follow the same pattern with slightly lower full-exposure ceilings: Marketing Managers (SOC 11-2021) intermediate exposure = 57.8%, full exposure = 93.8%; Market Research Analysts (SOC 13-1161.00) intermediate exposure = 57.7%, full exposure = 92.3%; Advertising and Promotions Managers (SOC 11-2011) intermediate exposure = 54.7%, full exposure = 83.7%. Public Relations Specialists, listed in the Eloundou paper as among the five most exposed occupations in the U.S. economy, score intermediate exposure = 78.8% — even higher than search marketing. We chose Search Marketing Strategist as the worked example rather than PR because matched field-experiment evidence (Hartmann et al., 2025; Lee et al., 2025) exists for AI-native marketing campaigns but not for AI-native PR work, allowing the productivity bracketing in 2.4 to be anchored to studies rather than projection. Marketing functions sit firmly in the high-exposure regime.

Figure 11: AI exposure scores for marketing occupations.

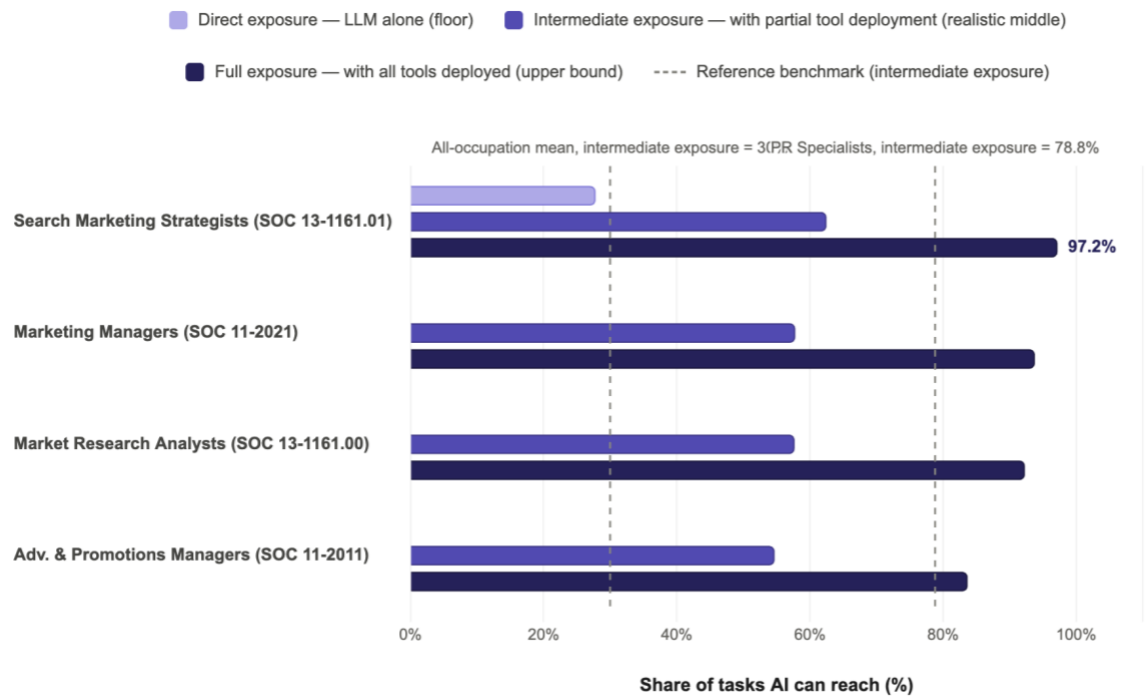

Note: This figure compares AI exposure levels across selected marketing occupations using the exposure framework developed by Eloundou et al. (2024). Direct exposure reflects tasks that current LLMs can perform independently, intermediate exposure reflects tasks reachable when AI is paired with complementary software tools, and full exposure represents the upper-bound scenario in which all available tools are fully integrated into workflows.

Grouped horizontal bar chart showing direct, intermediate, and full exposure for four marketing occupations. Search Marketing Strategist scores are this paper's task-level application of the rubric of Eloundou et al. (2024). (direct exposure 27.8%, intermediate 62.5%, full 97.2%; see 2.6). The other three occupations are shown for comparison. Two dashed reference lines mark Eloundou's all-occupation mean (intermediate exposure 30%) and Public Relations Specialists.

A full-exposure score of 97.2% means nearly every task in the occupation is technically reachable by AI under full tool deployment. Reachable, however, is not the same as should be handed over. Some tasks (brand positioning, regulatory judgment, creative strategy) are reachable by AI but should remain under human control because the consequences of getting them wrong are too high. That distinction is what the tier system handles: full exposure tells the size of the opportunity; the tiers tell what to do with it.

B.3 Derivation of the three productivity options

The published field-experiment evidence establishes a productivity gap between bolt-on AI deployment and deployment that restructures work around AI capability. This appendix documents the three reference points that bracket the gap. It does not derive a productivity prediction for OOA. OOA is offered as a defensible architecture for the redesign the gap implies; whether OOA-architected functions capture gains in any specific portion of the documented range is an empirical question this paper does not answer. The three options below describe deployment archetypes the literature has separately documented — Option 1 from Humlum and Vestergaard's 25,000-worker Danish field study, Option 2 from Fang et al.'s seven-workflow randomized experiments at a major consumer platform, and Option 3 from Hartmann et al. and Lee et al.'s field studies of AI-native creative production. Each option is a literature finding bracketing what AI deployment produces under different organizational conditions; the gap between Option 1 and Option 3 is what redesign addresses.

Option 1 — Bolt-on. Anchored to Humlum and Vestergaard (2025): 25,000 workers across eleven occupations in Denmark, ~3% time savings, no significant impact on earnings or hours. The baseline.

Option 2 — Integrated. Anchored to the workflow-level field-experiment range documented by Fang et al. (2025), who ran large-scale randomized experiments integrating generative AI into seven consumer-facing workflows at a cross-border online retail platform with millions of users. Treatment effects ranged from 0% to 16.3%, with the variation tracking how much the AI contributed relative to existing practice. This bracket — roughly 5–15% productivity gain at the workflow level — represents the realistic middle case where AI is integrated at the task level but the surrounding role architecture is largely intact. Per-task experimental results from Noy and Zhang (2023) on professional writing (40% time reduction, 18% quality increase at the per-task level under controlled conditions) are higher than the workflow-level Fang bracket because per-task gains do not compose multiplicatively to function-level productivity. The Fang range is the closer real-world analogue.

Option 3 — Restructured. Two field experiments anchor the upper bound. Hartmann et al. (2025) ran AI-generated banner ads against professional stock photography across more than 173,000 impressions; AI variants achieved up to 50% higher click-through rates at $0.04 per image versus $9 for stock and $100 for a commissioned freelancer. Lee et al. (2025) compared three creative approaches in a controlled experiment (human-created, AI-modified human, fully AI-created); the fully AI-created ads outperformed both alternatives by up to 19% in click-through rate, while AI-modified human ads showed no significant improvement.

B.4 Caveats and boundary conditions

Three methodological caveats apply.

First, the Eloundou rubric is one of several AI exposure measures in the literature. The exposure figures for search marketing come from applying that rubric — the most widely cited instrument — to the occupation's task list. Alternative measures broadly agree on which occupations are most exposed but disagree on exact magnitudes. The numbers should be read as central tendencies, not precise quantities.

Second, the Option 2 bracket is derived from field-experiment evidence on integrated AI deployment (Fang et al., 2025) rather than from extrapolation of per-task experimental results. Per-task results like Noy and Zhang's (2023) 40%-time reduction on professional writing are higher than the workflow-level Fang bracket because per-task gains do not compose multiplicatively to function-level productivity. The 5–15% Fang range is the closer real-world analogue.

Third, the Option 3 case rests on two field experiments drawn from one subset of marketing work — visual ad creative. A preregistered meta-analysis of 106 experiments and 370 effect sizes in Nature Human Behavior finds that human-AI combinations on average perform worse than the best of humans or AI alone (Hedges' g = −0.23), with content creation tasks as the main exception (Vaccaro et al., 2024). Marketing sits in the empirical sweet spot. The gains documented here do not generalize automatically to all knowledge work. Additional field experiments in other marketing task families — paid search optimization, attribution modeling, customer segmentation — would strengthen the base further.

These caveats locate the argument; they define the inferential weight the worked example can carry. The specific quantities the appendix produces — intermediate exposure of 62.5% for search marketing, the 5–15% Option 2 bracket, the 9-task Tier 1 / 18-task Tier 2 / 9-task Tier 3 distribution in the cross-tabulation — are illustrative. They show what framework application looks like, and they will vary by coder under the inter-rater patterns Yin et al. (2026) document. The framework's directional claims — that bolt-on deployment captures less than restructured deployment of comparable AI capability, and that Tier 1 work serves a developmental function whose preservation has implications for workforce supply — are robust to coder variation in the illustrative quantities. The worked example demonstrates how the framework is applied to a specific function so the directional claims can be tested concretely; it is not itself the test.

B.5 The 36 tasks of Search Marketing Strategist: full assignment to Eloundou category and OOA tier

This section provides the full task-by-task application of the Eloundou et al. (2024) exposure rubric and OOA tier-assignment criteria to the 36 tasks O*NET enumerates for Search Marketing Strategist (SOC 13-1161.01). The purpose is auditability: a typological theory in the Doty and Glick (1994) tradition has to be fully specified at the application level — independent researchers should be able to determine where each empirical case falls and disagree with specific assignments rather than guess at the authors' reasoning.

Sources and methodology. The task list is taken verbatim from the O*NET 13-1161.01 occupation summary (U.S. Department of Labor, retrieved April 2026). The Eloundou exposure category (E0/E1/E2) for each task is the authors' application of the rubric specified in Eloundou et al. (2024) — E0 indicates the rubric assessment that LLM use cannot reduce task-completion time by 50% at equivalent quality; E1 indicates LLM alone can; E2 indicates LLM with specialized software can. Eloundou et al. (2024) report economy-wide exposure means and distributions rather than scores for individual occupations; the task-level scores in this table are therefore the authors' application of the published rubric, not Eloundou et al.'s

own task-level scores. Independent application by other researchers using the same rubric may produce different assignments at the margin — recent work by Yin et al. (2026) finds substantial disagreement across frontier LLMs applying the rubric to the same O*NET task set, with mean exposure differing by a factor of up to 3.6 across models and pairwise agreement as low as 57%, suggesting that any single application should be read as one informed judgment rather than as a definitive scoring.

The OOA tier assignment for each task applies the criteria stated in section 2.2.3: how consequential the decision is, and what type of judgment it requires. Tier 1 captures low-consequence reversible high-volume work amenable to AI execution under light human supervision. Tier 2 captures medium-consequence work where Operator supervision decides when to trust, override, or escalate AI output. Tier 3 captures high-consequence strategic or expert-judgment work that stays under direct human authority.

Aggregate scores from this application. Applying the Eloundou rubric to the 36 tasks: 1 task scores E0, 10 score E1, 25 score E2. From these counts, the occupation's exposure measures are direct exposure (E1) = 27.8%, intermediate exposure (E1 + 0.5 × E2) = 62.5%, and full exposure (E1 + E2) = 97.2%. These figures are derived from the task-level application above rather than transcribed from an external source, so the three measures are arithmetically consistent by construction; they are the figures used in 2.2 and in the worked example. Eloundou et al. (2024) report economy-wide exposure means — roughly 14%, 30%, and 46% at the occupation level for the three measures — and do not publish a score for this specific occupation; search marketing's task-level scores sit well above those averages, consistent with Eloundou's finding that higher-wage, preparation-intensive occupations are the most exposed. This is one informed scoring under stated criteria, not a definitive measurement; the inter-rater variation that Yin et al. (2026) document applies.

The OOA tier distribution from this application: Tier 1 = 9 tasks (25%), Tier 2 = 18 tasks (50%), Tier 3 = 9 tasks (25%). The plurality concentration in Tier 2 confirms the worked example's framing of Search Marketing Strategist as an Operator-dominant occupation: the largest share of work sits in supervised-AI execution.

In the following table, each row gives the O*NET task identifier, the task text (verbatim from O*NET), the applied Eloundou category, the assigned OOA tier, and a one-line rationale for the tier assignment.

Table 1: Full task table based on Eloundou category for Search Marketing Strategist

O*NET ID	Task (abbreviated)	Eloundou	OOA Tier	Rationale
20328	Manage tracking and reporting of search activities; provide analyses to executives	E2	2	Reporting layer is procedural; analysis-for-executives requires Operator interpretation and stakeholder framing
20329	Optimize digital assets (text, graphics, multimedia) for	E1	1	Bulk SEO optimization is procedural, reversible, low individual consequence

	SEO and usability			
16192	Collect and analyze Web metrics (visits, CTR, conversion, CPA, CPC)	E2	2	Metric collection automatable; pattern interpretation requires Operator judgment about anomalies and brand context
16182	Participate in development or implementation of online marketing strategy	E2	3	Strategy development synthesizes brand, market, business context; Tier 3 Orchestrator work
16190	Optimize Web site exposure by analyzing search patterns to direct keyword placement	E1	1	Keyword research and placement at scale; AI can perform end-to-end
20320	Coordinate with developers to optimize Web architecture for search consumption	E2	2	Technical recommendations AI-draftable; cross-function tradeoffs require Operator judgment
20318	Assist in setting up or optimizing analytics tools for visitor tracking	E2	1	Configuration is procedural; once correctly set up, low ongoing consequence
20326	Identify KPIs and report key metrics from digital campaigns	E2	2	Reporting automatable; KPI selection requires judgment about what matters
20321	Create content strategies for digital media	E2	3	Content strategy is brand-shaping; sets editorial direction; Tier 3

20319	Combine secondary data with keyword research to profile user intent	E2	2	Data combination automatable; intent profiling requires Operator interpretation
16189	Collaborate with marketing staff to integrate strategies across sales channels	E0	3	Cross-functional negotiation and strategic integration; LLM cannot replace
20330	Optimize shopping cart experience or conversion rates against KPIs	E2	2	AI recommends; selecting and implementing requires Operator judgment about brand and revenue
20327	Improve search activities through analysis, experimentation, and A/B tests	E2	1	Test execution at scale is procedural and reversible; standard statistical inference
16191	Conduct online marketing initiatives (paid ads, affiliate, email, viral)	E2	1	Bulk paid placement and email at scale is procedural and reversible
16186	Conduct market research analysis (search trends, social topics, competitor performance)	E2	2	Data gathering automatable; identifying meaningful trends requires Operator judgment
16179	Propose online or multi-channel campaigns to marketing executives	E1	3	Drafts AI-assistable; executive proposal is high-consequence stakeholder work
20322	Evaluate emerging media or technologies and recommend application	E2	3	Strategic technology evaluation; expert judgment about fit with function and brand

16188	Communicate with merchants, webmasters, bloggers about hyperlink placement	E1	2	Outreach drafting AI-supportable; relationship management requires Operator judgment
16175	Identify, evaluate, or procure hardware or software for marketing campaigns	E2	2	Vendor evaluation AI-supportable; procurement requires judgment about budget, fit, integration
16193	Collaborate with design staff to create multimedia sites conforming to brand	E1	2	Content creation AI-augmented; brand-conformance review requires Operator judgment
16172	Keep abreast of regulations and emerging Web technology for compliance	E2	3	Compliance with regulation is high-consequence; expert judgment about applicability
20331	Purchase or negotiate placement of listings in local search engines and directories	E1	2	Bulk procurement automatable; negotiation requires Operator judgment
16185	Coordinate with merchandising/operations on product catalog accuracy and findability	E2	2	Cross-functional coordination requires judgment about competing priorities
20325	Execute or manage social media campaigns to inform search marketing	E1	1	Social campaign execution at scale is procedural, reversible

16187	Conduct financial modeling for online marketing programs and revenue forecasting	E2	2	Model construction AI-augmented; assumptions and forecasting choices require Operator judgment
16174	Implement online customer service processes for consistent user experience	E2	2	Process implementation AI-supportable; experience tradeoffs require Operator judgment
16184	Develop transactional Web applications using HTML/XML	E1	1	Code generation for HTML/XML is now an LLM strength; standard transactional patterns procedural
16183	Identify commercial/technical specs (usability, pricing, checkout, data security)	E2	3	Pricing and data security specifications are high-consequence; expert judgment
16177	Define product requirements based on market research with UI design and engineering	E2	3	Product requirement definition is high-consequence design work; cross-functional
20324	Execute or manage banner, video, or non-text link ad campaigns	E2	1	Banner/video ads at scale (Hartmann et al. 2025; Lee et al. 2025): AI-native production empirically outperforms; Tier 1 with Operator oversight
16178	Assist in evaluation or negotiation of contracts with vendors or online partners	E2	2	Contract analysis AI-supportable; negotiation requires Operator judgment about terms and relationship

16181	Prepare e-commerce designs or prototypes (storyboards, mock-ups)	E1	1	Prototype generation now an LLM/diffusion-model strength; reversible, iterative
16180	Assist in development of online transaction or security policies	E2	3	Security policy is high-consequence; legal and risk implications
20323	Execute and manage communications with digital journalists or bloggers	E1	2	Outreach drafting AI-supportable; press judgment requires Operator review for brand exposure
16173	Resolve product availability problems with customer service staff	E2	2	Cross-function problem-solving requires Operator judgment about customer-facing tradeoffs
16176	Identify methods for interfacing Web applications with ERP or other systems	E2	2	Integration design is technical-architecture work; AI-supportable but requires Operator judgment about system tradeoffs

The Eloundou × Tier cross-tabulation reveals the layered structure that motivates the framework.

Table 2.2. The cross-tabulation for Tier and Eloundou category

	Tier 1 (autonomous AI)	**Tier 2 (Operator-supervised)**	**Tier 3 (human authority)**	**Row total**
E0 (no meaningful AI capability)	0	0	1	1
E1 (LLM alone)	5	4	1	10

E2 (LLM + specialized software)	4	14	7	25
Column total	**9**	**18**	**9**	**36**

The two layers (capability and governance) separate cleanly: Eloundou tells us what AI can technically perform; OOA tells us where human authority is required regardless. Tier 3 contains 8 tasks that AI can technically support (E1 or E2) but that the framework reserves for human authority on consequentiality grounds — exactly the conceptual distinction the section 2.3 mapping discussion develops.

Limitations of this audit. This application is one researcher team's judgment under stated criteria. Three known sources of variation should be expected on independent replication: (a) the Eloundou rubric is sensitive to the model used for application — Yin et al. (2026) document a 3.6-fold divergence in mean exposure across ChatGPT-5, Gemini 2.5, and Claude 4.5 applying the same rubric to the O*NET task universe, with pairwise agreement as low as 57%; (b) the consequentiality criterion involves judgment that varies by industry context — what is high-consequence in regulated finance may be medium-consequence in growth-stage startups; (c) the Tier 1/Tier 2 boundary is fuzziest — specifically, whether tasks like analytics setup or A/B testing belong to autonomous AI execution or supervised-AI execution depends on assumptions about how much oversight is required for the specific deployment context. Independent coders applying the same criteria to the same task list should converge at the Tier 3 boundary (high-consequence work) more than at the Tier 1/Tier 2 boundary; this prediction itself is testable.

Appendix C: The OOA Framework Applied, Two Contrasting Functions

The search marketing example in section 2.2.3 showed the framework applied to an Operator-dominant function. The biggest gains sat in Tier 2, where supervised execution turned AI output into campaign results. The two examples that follow test the framework against very different functions: healthcare clinical decision support, which is Artisan-anchored, and credit underwriting, which is Orchestrator-anchored. One point to clarify before those examples begin. To say a function is “anchored” in a tier is to say that the function's primary risk lives there. Primary risk means the point at which a wrong decision produces consequences that cannot be undone by later review. This is the criterion used to place each function. In marketing, primary risk is commercial and reputational. It sits in Tier 2, where humans decide which AI-generated campaigns actually go live. In healthcare, primary risk is clinical harm to the patient. It sits in Tier 3, where no algorithm substitutes for the attending physician's judgment. In credit, primary risk is systemic and regulatory — model failure, fair-lending violation, portfolio miscalibration. It sits in Tier 3, where the Orchestrator decides what the automation is permitted to do. The same criterion applies in all three cases. The archetype distributions differ because the functions govern different kinds of decisions.

Function 1: Clinical Decision Support — An Artisan-Anchored Function

A physician's work resists the clean task-decomposition that made the search marketing example tractable. The O*NET inventory for General Internal Medicine Physicians (SOC 29-1216.00) lists

nineteen tasks. They span documentation, diagnostic reasoning, direct patient care, and clinical communication. But the tasks are not independent units that can be scored and redistributed one at a time.

A physician diagnosing a patient with multiple overlapping conditions does several things at once. She performs the physical exam, interprets test results, weighs treatment risks against patient preferences, and communicates prognosis. The quality of that encounter depends on doing these activities together — not on doing each one well in isolation.

This is the structural reason healthcare sits lower on the Eloundou exposure spectrum than text-heavy knowledge-work functions. The Eloundou paper itself finds that science and critical thinking skills are strongly negatively correlated with AI exposure. Tasks that require integrated judgment cannot be broken into separate pieces an LLM accelerates on its own (Eloundou et al., 2024).

AI is still transforming the function — but along a distribution that concentrates automation in Tier 1 while leaving Tier 3 human control over clinical decisions intact. Microsoft's Nuance DAX Copilot illustrates the Tier 1 layer. The system listens to patient encounters in real time, drafts structured clinical notes, and presents them for physician review. A Microsoft-administered July 2024 survey of 879 clinicians across 340 healthcare organizations reported an average of five minutes saved per clinician per encounter (Microsoft, 2024).

Liu, Hetherington, Dharod, et al., 2024The independent evidence is more tempered. A longitudinal study in NEJM AI tracked 112 primary care clinicians using DAX against 103 controls over 180 days at Atrium Health. The primary analysis found no statistically significant improvement in EHR or financial metrics. Gains appeared only in three exploratory subgroups — clinicians with the highest DAX adoption rates, low-volume clinicians, and family medicine clinicians (Liu, Hetherington, Dharod, et al., 2024). The pattern matches what section 2.2.1 documented in the productivity literature. The technology produces measurable gains in controlled settings. But the gains depend on workflow redesign, specialty fit, and clinician adoption — they are not automatic, and they do not distribute evenly across users.

The Tier 2 and Tier 3 layers are where the framework's predictions about human-AI collaboration get tested most directly. The current evidence is cautionary rather than conclusive.

A randomized clinical trial in JAMA Network Open assigned fifty physicians (26 attendings and 24 residents) to diagnose clinical cases with or without access to GPT-4. The LLM group scored 76% on diagnostic reasoning; the control group scored 74%. The adjusted difference was not statistically significant (95% CI, −4 to 8; P = .60). But GPT-4 alone scored 16 points higher than the control group (95% CI, 2 to 30; P = .03) (Goh et al., 2024). The result establishes the empirical signature of the Tier 2 Operator problem: the AI alone produces superior outputs on specific diagnostic tasks, yet the combined system of physician-plus-AI does not systematically outperform physicians using conventional resources. The gap between what the AI can do on its own and what the human extracts from it is exactly what Tier 2 governance design has to address. A subsequent NEJM AI Script Concordance Testing benchmark across 750 questions and ten international datasets found the best-performing LLM (OpenAI's o3, at 67.8%) still did not reach the level of senior residents or attending physicians on decisions under uncertainty (McCoy, Swamy, Sagar, et al., 2025). The Goh finding may not generalize to all clinical AI, particularly structured decision support rather than open-ended diagnostic reasoning.

Within the governance architecture, the redistribution is therefore Artisan-anchored under the risk-locus criterion. Primary risk in clinical decision support is the risk of patient harm. That risk sits in Tier 3, where the attending physician exercises judgment no algorithm yet substitutes for.

The Tier 1 automation target is larger than most leaders realize. The AMA's 2024 national physician survey found that of the average 57.8-hour physician workweek, only 27.2 hours — less than half — go to direct patient care. Another 13 hours go to indirect work: documentation, order entry, and interpreting test results (AMA, 2025). That 13-hour block is what AI ambient scribes, automated order entry, and clinical decision support are absorbing first. Tier 2 Operator functions cover AI-generated diagnostic suggestions, imaging flags, and sepsis alerts — outputs that need clinical review before action.

The irreducible value — the work that defines what a physician actually is — concentrates in Tier 3. This is complex diagnostic reasoning under genuine uncertainty. Therapeutic relationships built on trust and empathy. Ethical decisions about end-of-life care. Physical examination, which no algorithm replicates. The FDA's regulatory architecture enshrines this distribution. A systematic review in JAMA Network Open found that of 950 AI-enabled medical devices authorized through June 2024 — 76% of which were in radiology — 97% were cleared via the 510(k) pathway under the explicit assumption of a human-in-the-loop (Sivakumar et al., 2025). The FDA's own running database has since expanded that total to approximately 1,451 authorized devices by the end of 2025 (FDA, 2025). The Artisan — the attending physician whose expertise makes AI outputs evaluable — is not optional. The regulatory structure presumes the Artisan exists.

The pipeline threat is correspondingly severe. A 2026 Perspective in npj Digital Medicine argues that clinical AI should be treated as a "digital copilot," not an "autopilot." The authors — drawing on collaboration between UCL/Moorfields Eye Hospital clinicians and Lufthansa flight safety experts — warn of two distinct risks they call never skilling and mis-skilling. These differ from deskilling. Deskilling is the gradual erosion of existing competence. Never skilling is the failure to develop foundational skills in the first place. The authors draw a direct parallel to aviation's "children of the magenta line" — pilots so dependent on automated flight systems that they could not fly manually when the automation failed (Ong, Merle, Pollreisz, et al., 2026).

In medicine, the parallel is a generation of residents who learn to review AI-generated notes rather than write them, accept AI-flagged diagnoses rather than reason through differentials independently, and rely on AI-scored imaging rather than develop their own pattern recognition through thousands of hours of deliberate practice. The Artisan's capability is built through exactly the repetitive, effortful, consequential work that AI is now absorbing from the training pipeline. Unlike in marketing or finance, the consequences of a hollowed Artisan pipeline in medicine are measured in patient lives.

Function 2: Credit Underwriting — An Orchestrator-Anchored Function

Credit underwriting sits at the opposite end of the OOA spectrum from healthcare. Under the risk-locus criterion, the Orchestrator is primary. The decisive risks — model failure, fair-lending violation, portfolio miscalibration — do not concentrate in any individual loan decision. They concentrate in the system architecture that produces millions of loan decisions at once.

The O*NET inventory for Loan Officers (SOC 13-2072.00) lists thirty tasks. They are dominated by data analysis, financial statement interpretation, risk assessment, report writing, and payment computation — almost all text- and number-heavy operations that fall squarely within the LLM exposure frontier. Credit Analysts (SOC 13-2041.00) carry an even more concentrated profile: eleven tasks, nearly every one classifiable as E1 or E2 under the Eloundou rubric. This makes credit analysis one of the most AI-exposed occupations in financial services (Eloundou et al., 2024).

Where healthcare splits between high-exposure documentation and low-exposure embodied clinical work, credit underwriting's inventory is uniformly exposed. The question is not whether AI can accelerate these

tasks. It is how much of the function can be fully automated before human governance becomes the binding constraint.

The deployment data makes the scale visible. Upstart's AI lending platform serves more than 100 lending partners as of its 2024 annual report, evaluates applicants through more than 2,500 variables that extend well beyond the traditional FICO score, and automates more than 91% of loan originations with no human involvement (Upstart, 2024 10-K). Zest AI serves nearly 300 lender customers with more than 650 proprietary credit models. Its partner institutions report approval increases of 25% with no added risk, and default reductions of 20% at constant approval rates. One partner, Commonwealth Credit Union, reports AI decisions on 70–83% of consumer loan applications (Zest AI, 2025; Zest AI, 2024).

A study in MIS Quarterly examined a major bank with over fifty million customers adopting AI credit scoring for underserved populations. The AI model achieved higher prediction accuracy than traditional approaches overall — but with one important heterogeneity. Applicants with the thinnest credit files, those with what the authors describe as "missing weak signals," saw smaller improvements in approval rates than applicants with more signal-rich records (Li et al., 2024). AI expanded access to credit broadly but did not solve the hardest problem. Borrowers with the sparsest data remained the hardest to score accurately. This sharpens the Orchestrator's actual job. It is not merely governing model quality in the aggregate. It is designing systems that detect and correct where AI systematically underperforms for the populations financial inclusion policy most wants to reach.

One boundary condition deserves mention. These automation rates are drawn from consumer personal lending — a high-volume, relatively standardized product. Commercial credit underwriting is different. Deal structures are bespoke, borrower relationships are complex, and the function retains a substantially larger Artisan component. The automation frontier there is materially lower.

This is Tier 1 automation operating at a scale that fundamentally redefines the function. When nine out of ten loan applications are decided by an algorithm with no human involvement, the Operator layer shrinks to a narrow band of the workflow. Primary human value migrates upward, to the Orchestrator who designs the system architecture. Which variables does the model include? What thresholds trigger escalation to human review? How is fair-lending compliance engineered into the decisioning pipeline? When does the portfolio-level output warrant override? The Brookings Institution captures this redistribution precisely: in AI-augmented financial services, the most valuable professionals are not those with perfect models but those who know when not to trust them (Levy Yeyati, 2025).

The regulatory architecture makes the Orchestrator role not merely valuable but legally mandatory. The Federal Reserve's SR 11-7 model risk management guidance requires independent validation, ongoing monitoring, and comprehensive documentation of every quantitative model a bank deploys (Board of Governors of the Federal Reserve System, 2011). The scale of this mandate is substantial. The Risk Management Association's 2024 survey found the average bank maintains 175 quantitative models in its inventory. An earlier RMA survey (2021/2022) documented a range from single digits at community banks to 800 at the largest institutions (RMA, 2024).

The CFPB requires creditors to provide specific, accurate reasons for credit denials even when using complex algorithms. Creditors cannot cite algorithmic opacity as a basis for failing to explain adverse decisions (Consumer Financial Protection Bureau, 2023). The EU AI Act (Regulation 2024/1689) entered into force on August 1, 2024; its high-risk provisions become fully enforceable on August 2, 2026. It classifies credit scoring explicitly as high-risk under Annex III section 5(b), imposing mandatory human oversight design requirements, risk management systems, and transparency obligations (European Parliament, 2024). Colorado's SB 24-205, one of the first comprehensive U.S. state-level AI laws, had its effective date amended from February 2026 to June 30, 2026. It requires annual impact assessments and

grants consumers the right to appeal AI-made credit decisions through human review (Colorado General Assembly, 2024).

Collectively, these frameworks create a Tier 3 governance layer that is structurally inseparable from the function itself. An Orchestrator who understands model architecture, regulatory constraints, and fairness testing is not an optional enhancement to the credit function. The Orchestrator is the condition under which the automation is legally permitted to operate.

The pipeline threat in credit underwriting has the same structural logic as in healthcare, but it operates through a different mechanism. A 2026 analysis by the Federal Reserve Bank of Dallas, introduced in section 2.2.1, argues that AI substitutes for entry-level workers but complements experienced ones. Codifiable knowledge is automatable. Tacit knowledge — the intuitive risk judgment that compounds with thousands of applications reviewed — is complementary to AI. Insurance underwriters, credit analysts, lawyers, and marketing specialists show experience premiums exceeding 100%. Senior professionals earn more than double their entry-level counterparts, precisely because tacit judgment accumulates through repetition (Davis, 2026).

But when Upstart automates more than 91% of loan originations and Zest AI partner institutions auto-decision 70–83% of consumer loans, the traditional developmental pathway breaks down. Junior underwriters used to learn credit assessment through high-volume manual review, building the pattern recognition that eventually qualified them for complex commercial decisions. That volume is gone. The applications that still reach a human are precisely the edge cases too complex for beginners. An Orchestrator-tier role under OOA would require professionals who understand both the domain and the system architecture governing its behavior. The framework's pipeline-preservation prescription specifies that a defined share of Tier 1 work be reserved as developmental exposure precisely because that understanding has historically been built through years of performing the routine work the AI has now absorbed. Whether the prescription, applied to a credit function, would in fact preserve the pipeline is an empirical question that field deployment over a 3–5-year window would test; the architectural specification is the framework's contribution. As the Dallas Fed concludes, leaving new employees off the job ladder is not sustainable in the long run.

Synthesis: What the Two Examples Demonstrate

Read alongside search marketing, these two examples establish four findings.

First, the OOA archetype distribution is domain-conditional. Marketing concentrates human value in the Operator layer. Healthcare concentrates it in the Artisan layer. Credit concentrates it in the Orchestrator layer. Any organization applying a uniform AI augmentation strategy across functions will misallocate investment at every point where this distribution differs.

Second, the anchoring tier in each function is determined by the locus of primary risk. In marketing, commercial and reputational risk concentrates in the supervised-execution layer. In healthcare, clinical risk to patients concentrates in the Tier 3 judgment layer. In credit, systemic and regulatory risk concentrates in the Tier 3 system-design layer. The three functions are anchored in three different tiers under a single consistent criterion — where decisive wrongness lives. This is the strongest generalization the three examples jointly earn, and it gives the framework a principled way to predict the anchoring tier for functions not yet examined.

Third, the regulatory environment determines the minimum viable governance architecture. FDA 510(k) clearance presumes a clinician-in-the-loop. SR 11-7 and the EU AI Act mandate independent model

validation. In regulated industries, governance tiers are legal requirements, not organizational design choices.

Fourth, the pipeline threat operates through the same mechanism across all three domains — the automation of developmental work — but its severity escalates with the irreversibility of the decisions the function governs. In marketing, a hollowed pipeline produces campaigns that lack strategic coherence. In credit, it produces model governance that cannot detect when the models are wrong. In healthcare, it produces clinicians who cannot function when the AI fails.

A final methodological note. These two examples illustrate how the framework applies across functions with different archetype distributions. They do not constitute a formal test of OOA in the falsificationist sense, and the appendix should not be read as establishing that the framework has been validated against the possibility of failure. What they establish, more modestly, is that the framework's architecture is consistent across domains with very different task structures, regulatory environments, and risk loci — the same two-input assignment rule produces coherent redistributions in marketing (Operator-anchored), healthcare (Artisan-anchored), and credit (Orchestrator-anchored). The architecture is consistent across domains; whether application of the architecture delivers the productivity outcomes the field-experiment literature documents is a separate empirical question this paper does not answer. That cross-domain consistency of architecture is the minimum condition under which an integrated test becomes worth designing.

An organization that decomposes its own functions through the same method — mapping tasks to exposure levels, routing them through governance tiers, and identifying which archetype anchors its irreducible human value — will arrive at a workforce architecture specific to its context rather than borrowed from a generic playbook. That is the framework's intended use.

Appendix D: This Paper's Framework Among Other Frameworks

This appendix asks where this paper's framework sits among the other frameworks for AI workforce transformation now in circulation.

As section 2.2.3 sets out, the framework's contribution is architectural: a function-level specification of what AI workforce redesign should consist of, organized as two prescriptions — a redistribution prescription and a pipeline-preservation prescription. This appendix tests the closest competing frameworks against both. It scores the practitioner frameworks on the two prescriptions, examines the three engaged in depth, and applies each — alongside this framework — to the Search Marketing Strategist worked example from section 2.2.3.

D.1 The two-prescription gap matrix

To compare frameworks fairly, the appendix scores each on both prescriptions using explicit criteria.

On the redistribution prescription, frameworks are scored by their scope of prescription. None if the framework does not propose a method for redistributing work between humans and AI systems. Diagnostic if it identifies the redistribution problem and provides survey or benchmarking data without prescribing a method. Function-level prescriptive if it prescribes a redistribution at the level of an individual occupation or function, leaving enterprise-level operating model and topology outside its scope. Enterprise-level prescriptive if it prescribes restructuring at the level of the operating model,

organizational topology, or business model — agentic teams, total reinvention, work-as-product redesign across the enterprise. The Function-level and Enterprise-level categories are equally substantial in change demanded; they differ in the unit they restructure, not in how heavy a lift each requires. Whether function-level scope is empirically superior to enterprise-level scope, or vice versa, is an open question — no head-to-head causal study currently adjudicates the comparison.

On pipeline preservation: None if the framework's main publications do not engage with apprenticeship, entry-level disruption, or how novices become experts. Adjacent if it engages with related concerns (changing skills, reskilling, role evolution) without addressing how the developmental pipeline gets preserved. Acknowledged if it names the pipeline problem in its main publications but does not propose structural solutions. Substantive if it proposes specific organizational design responses for keeping the pipeline intact.

By these criteria, the practitioner frameworks engaged in depth in this appendix span the closest competitors at three distinct units of analysis. McKinsey's Agentic Organization is the closest competitor at the enterprise-architecture level. Deloitte's 2025 Global Human Capital Trends with the Humans × Machines work-design framework is the closest competitor at the human-capital-benchmarking-and-redesign level. Daugherty and Wilson's Human + Machine is the closest competitor at the worker-role-vocabulary level. Each operates at a different scope and answers a different question; together they cover the practitioner landscape's closest engagements with the workforce-design problem this paper addresses.

Table 2: This paper's framework against the closest practitioner competitors

Framework	Source	What it tells an organization to do	Productivity & governance	Pipeline preservation	Key limitation

McKinsey *Agentic Organization*	Sukharevsky et al., Sept 2025	Dissolve traditional functions into agentic teams of 2–5 humans supervising 50–100 AI agents. Staff with M-shaped supervisors and T-shaped experts. Add critic, guardrail, and compliance agents for governance. The most architecturally complete enterprise-level prescription available.	**Enterprise-level prescriptive**	**None** (assumes M-shaped supervisors already exist; does not explain how they get developed when AI absorbs the entry-level work that historically produced them)	Best answer to "what does the AI-native enterprise look like?" Silent on how the next generation of supervisors gets built. The 2-to-5-humans-per-50-to-100-agents ratio is drawn from consulting engagements rather than peer-reviewed empirical evidence; whether the ratio is stable across functions has not been independently tested.

Deloitte *2025 Global Human Capital Trends / Humans × Machines*	Mallon, Cantrell, & Forsythe, Mar 2025	Surveys ~10,000 leaders across 93 countries. Finds 66% of managers and executives say recent hires are not fully prepared, experience is the #1 gap, and 57% of executives gained initial experience through apprenticeships or entry-level roles. Reports that only 16% of organizations have fully redesigned roles, processes, and operating models around AI. Proposes Work Engineering (work as a product) plus bridging solutions: skills-based hiring, alternative credentials, in-context upskilling.	**Diagnostic** with **Enterprise-level prescriptive** elements (Work Engineering implies full redesign but the public methodology does not provide task-level specifications)	**Acknowledged** (names the pipeline problem with real data; proposed responses bridge the gap rather than preserve the developmental work itself)	Best diagnosis of the pipeline problem in the practitioner space. The bridge-vs-preserve distinction is where this paper's contribution begins. Public framework does not specify which tasks to redistribute or where tier boundaries should fall.

Daugherty & Wilson *Human + Machine*	HBR Press 2018, updated 2024	Maps the "missing middle" between pure-human and pure-machine work. Defines three roles (Trainers, Explainers, Sustainers) and three fusion skills (intelligent interrogation, judgment integration, reciprocal apprenticing). Best vocabulary for individual-level collaboration.	**Diagnostic** (does not distinguish governance intensity across tasks or predict gains from redistribution)	**Adjacent** (reciprocal apprenticing describes mutual human-AI learning, not how the workforce pipeline stays intact)	Strongest role vocabulary in the field. Does not score tasks, does not set tier boundaries, does not address how Trainers, Explainers, and Sustainers get developed in the first place.

This paper's framework (OOA + M-shape) *(Rashidi et al.)*	This paper, 2026	Score every task in an occupation against the Eloundou exposure rubric. Route each to one of three governance tiers: Tier 1 (automate), Tier 2 (Operator supervises AI), Tier 3 (Orchestrator/ Artisan controls). Reserve a defined share of Tier 1 tasks for new-entrant development to keep the expertise pipeline intact.	**Function-level prescriptive** (task-level redistribution at the level of an individual occupation; substantial restructuring within the function but does not prescribe enterprise topology, operating model, or business model)	**Substantive** (explicitly reserves developmental tasks for new entrants, even when AI could perform them faster, because those tasks build the expertise Tier 2 and 3 depend on)	Function-by-function deployment scope, not a whole-enterprise architecture. No field deployment data yet. Relies on a single exposure rubric (Eloundou) that alternative instruments may score differently. Three-tier structure is a taxonomic choice rather than an empirical finding. Designed to be embedded within broader transformation efforts, not to replace them.

D.2 The three practitioner frameworks engaged in depth

McKinsey Agentic Organization

McKinsey's "The Agentic Organization," published September 2025, is the most architecturally complete framework in the current practitioner landscape. The eight authors — Sukharevsky, Krivkovich, Gast, Storozhev, Maor, Mahadevan, Hämäläinen, and Durth — propose a model that operates across five enterprise pillars: business model; operating model; governance; workforce, people, and culture; and technology and data (Sukharevsky et al., 2025).

The basic organizational unit is the agentic team — a small group of multidisciplinary humans who own and supervise AI workflows that deliver end-to-end business outcomes. McKinsey's central structural claim, drawn from consulting engagements, is that two to five humans can supervise an agent factory of 50 to 100 specialized agents running a complete process such as customer onboarding, product launch, or month-end close. Three role archetypes populate these teams: M-shaped supervisors (broad generalists fluent in AI), T-shaped experts (deep specialists who handle exceptions), and AI-augmented frontline workers. The most distinctive contribution is agent-on-agent governance: critic agents challenge outputs, guardrail agents enforce policy, and compliance agents monitor regulation. The 2-to-5-humans-per-50-to-100-agents ratio is drawn from McKinsey's consulting practice rather than from empirical research; whether the ratio is stable across functions, industries, and regulatory contexts has not been independently tested.

On the redistribution prescription, McKinsey's framework is Enterprise-level prescriptive. It prescribes restructuring at the level of operating model, organizational topology, and governance: agentic teams, agent-on-agent governance, M-shaped supervisors, T-shaped specialists. The scope is the whole enterprise, and the unit of restructuring is the operating model. This paper's framework operates at a different unit — the individual occupation — and the choice between them, or whether to use both with this paper's framework embedded inside an enterprise-level transformation, depends on the organizational question being asked.

On the pipeline-preservation prescription, the framework is silent in its primary publications. The M-shaped supervisor and T-shaped expert presuppose that the relevant experts already exist. The framework offers no account of how those experts continue to be produced when AI absorbs the work that historically produced them. For an organization building the full enterprise of the AI era, McKinsey's framework is the more comprehensive starting point. For an organization that wants occupation-level architectural specificity and a sustainable supply of qualified Operators and Artisans, this paper's framework is the more directly responsive complement.

Deloitte 2025 Global Human Capital Trends and Humans × Machines

Deloitte's 2025 Global Human Capital Trends, titled Turning Tensions into Triumphs, was published in March 2025 under the leadership of Sue Cantrell and David Mallon. The "Closing the Experience Gap" chapter is authored by David Mallon, Sue Cantrell, and John Forsythe. The report draws on a survey of nearly 10,000 business and HR leaders across 93 countries, plus extra worker, manager, and executive surveys totaling more than 13,000 respondents. Its empirical reach on AI workforce dynamics is unmatched in the practitioner space. A separate Humans × Machines service offering layers on top, from which Deloitte reports that only 16% of organizations have fully redesigned roles, processes, and operating models around AI, and that companies taking a tech-first approach are 1.6 times more likely to report their AI investments are not exceeding expectations (Deloitte, 2025).

On the redistribution prescription, Deloitte's framework is Diagnostic rather than prescriptive at the architectural level. The 16% finding is a powerful organizational benchmark. The Humans × Machines Work Engineering service offering — which treats work as a product to be designed, tested, and continuously improved — is Enterprise-level prescriptive in implication, but the public methodology is not specified at a level that allows independent application. An organization that uses Deloitte's diagnostic learns that it is behind but does not learn from the public framework which specific tasks within which specific occupations to redistribute, or where to set the boundaries between governance tiers.

On the pipeline-preservation prescription, Deloitte's framework is Acknowledged, and its “Closing the Experience Gap” chapter is the closest practitioner analog to this paper's central concern. The chapter argues that as AI absorbs entry-level work, apprenticeship models erode, and that “the number of

traditional entry-level roles is shrinking, making it harder for workers to gain the necessary hands-on experience." It documents the resulting bind: "most so-called entry-level jobs now require two to five years of experience." It quotes Peter Cappelli capturing the situation: "Everybody wants to hire somebody with three years' experience, and nobody wants to give them three years' experience." The framing is honest. The data is solid. The diagnosis is right.

The substantive difference between Deloitte and this paper's framework lies in how each framework responds. Deloitte treats the experience gap as a hiring problem and proposes bridging solutions: skills-based hiring, alternative credentials, AI-mediated learning environments. This framework treats the same gap as an expertise-formation problem and proposes an architectural response: explicitly reserve a defined share of Tier 1 tasks within an occupation as developmental exposure for new entrants, even when AI could perform those tasks faster. The two responses are complementary. The diagnostic tells you the pipeline is breaking; the architecture tells you which specific tasks to preserve, in which specific occupations, to fix it.

Daugherty & Wilson Human + Machine

Daugherty and Wilson's Human + Machine (HBR Press 2018, expanded edition 2024) is the most durable framework in the practitioner landscape on what individual workers do in AI-augmented work. Its central concept is the missing middle — the zone of human-AI collaboration that sits between purely human work and purely machine work. Within that zone, the original 2018 framework defined three new role types: Trainers (humans who teach AI systems how to perform), Explainers (humans who interpret AI behavior for stakeholders), and Sustainers (humans who maintain and govern AI systems in production). The 2024 HBR update added three fusion skills that workers in the missing middle need: intelligent interrogation, judgment integration, and reciprocal apprenticing (Wilson & Daugherty, 2024).

Daugherty and Wilson's framework is the closest practitioner competitor to OOA at the worker-role level. Both frameworks operate at the level of what individuals do in AI-augmented work, and both propose three categorical roles. But the two role typologies are answering different questions. Trainers, Explainers, and Sustainers describe what kinds of work humans do alongside AI systems — they are task-types, not governance tiers. A worker can simultaneously be a Trainer (teaching the AI), an Explainer (interpreting its outputs to clients), and a Sustainer (catching errors in production); the three are not mutually exclusive role assignments. OOA's three archetypes, by contrast, partition workers by governance authority over AI outputs. An Orchestrator designs the workflow boundary; an Operator supervises AI execution above the loop; an Artisan provides the irreducibly expert audit. The three are mutually exclusive at the level of any given task: a task is either routed to autonomous AI execution, supervised by an Operator, or controlled by a human exercising Tier 3 authority. The two frameworks are complementary rather than competing — Daugherty and Wilson's vocabulary describes the texture of human-AI collaboration at the task level; OOA prescribes the governance structure at the role level.

On the redistribution prescription, Daugherty and Wilson's framework provides the best available vocabulary for what individual contributors do when they work well with AI. But it does not distinguish governance intensity across tasks: all missing-middle work is treated as collaborative without specifying which collaboration requires what level of human control. The role taxonomy is not designed to predict the productivity gain from redistributing tasks across tiers. On the pipeline-preservation prescription, the role taxonomy assumes the existence of analysts capable of training, explaining, and sustaining without theorizing how those analysts come to be. Reciprocal apprenticing, added in the 2024 update, describes mutual learning between an individual worker and an AI system, not how the workforce-level apprenticeship pipeline stays intact across cohorts. The metaphor is useful in its own right but it is not an answer to the pipeline-preservation question this paper engages.

D.3 A side-by-side worked example: Search Marketing Strategists

The fairest comparison of competing frameworks is to apply each to the same problem and see what each prescribes. section 2.2.3 develops Search Marketing Strategists (O*NET SOC 13-1161.01) as a worked example of OOA in action. Under the Eloundou rubric, nearly every task in the occupation is potentially exposed once LLMs are deployed within complementary software — a full-exposure score of 97.2% in this paper's task-level application (see 2.6) Field experiments demonstrate large productivity gains under the right conditions. The four prescriptions below show what each framework would say. The differences are revealing on both prescriptions — redistribution and pipeline preservation.

What this paper's framework (OOA + M-shape) prescribes

Decompose the O*NET task inventory and score each task with the Eloundou rubric. Allocate tasks to three governance tiers. Tier 1 (full automation): SEM bid management, A/B test execution at scale, keyword research and clustering, bulk creative variation generation, standard performance reporting. Tier 2 (Operator — human supervision of AI execution): campaign performance interpretation, anomaly investigation, competitive analysis, creative brief evaluation, exception handling for high-spend or brand-sensitive campaigns. Tier 3 (Artisan/Orchestrator — human-controlled work): channel strategy, brand positioning, attribution model design, complex client diagnosis, creative strategy.

On productivity, the redistribution this framework prescribes is positioned within the documented gap between bolt-on deployment (~3% time savings under no redesign) and the field-experiment ceiling (19–50% performance gains under production-layer restructuring). Whether OOA-architected functions capture gains within that gap, and where in the gap, is not established by the cited literature. What the framework establishes is the architectural specification: task-level scoring against the Eloundou rubric, three-tier governance assignment under stated criteria, role redesign into Orchestrator, Operator, and Artisan archetypes, and the explicit reservation of developmental tasks. The redistribution is substantial, prescribed at the level of the marketing function, and does not dissolve marketing into agentic teams or specify the enterprise topology of which marketing is part.

On the pipeline-preservation prescription, OOA adds one architectural specification no other framework makes: a defined share of Tier 1 tasks must be preserved as developmental exposure for new analysts even when AI could perform them faster, because that exposure is the substrate from which Tier 2 judgment gets built. An Operator who has never personally executed bid management, never investigated an anomaly without AI assistance, never written a creative brief from scratch cannot evaluate when the AI version of those tasks has gotten it wrong. The reservation is explicit, defined, and small enough to preserve nearly all the productivity gains while keeping the developmental pathway intact.

What McKinsey Agentic Organization would prescribe

Reorganize the search marketing function as an agentic team. Two to five human marketing professionals supervise an agent factory of 50 to 100 specialized agents — a bid-management agent, a creative-variation agent, a performance-reporting agent, an anomaly-detection agent, a competitive-monitoring agent, and so on. An M-shaped supervisor orchestrates the factory. T-shaped specialists handle exceptions. AI-augmented frontline workers spend more time with clients and less on systems. Agent governance is implemented through critic, guardrail, and compliance agents.

On productivity, McKinsey prescribes at the enterprise level: the marketing function dissolves into agentic teams within a broader operating-model redesign, and the prescription is integrated by design with restructuring across other functions and across governance. This paper's framework prescribes at the level

of the marketing function alone: a redistribution within marketing that does not depend on the operating model around it changing in lockstep. Both prescriptions require substantial restructuring within their respective scope. On pipeline, McKinsey's framework is silent. It does not specify how the M-shaped supervisor or T-shaped specialist gets developed, which earlier-career exposure produced them, or how the next generation will reach equivalent capability when all the developmental work is now handled by the agent factory. The two-to-five-humans-per-fifty-to-a-hundred-agents ratio is itself a pipeline-collapse risk: if every marketing function in the economy adopts it, the inflow of humans capable of supervising the factory falls below replacement.

What Deloitte Humans × Machines would prescribe

Start with diagnosis. Has the search marketing function actually redesigned roles, processes, and operating models around AI integration, or has AI been layered onto existing structures? If the function falls into the 84% of organizations that have not redesigned, apply Work Engineering: treat the work as a product, with intentional human-machine task allocation, continuous testing, and iterative improvement. Use the experience-gap diagnostic to monitor whether mid-level talent is becoming harder to develop or hire. When the pipeline shows weakness, deploy bridging solutions: skills-based hiring criteria, alternative credentials such as portfolio assessments and structured simulations, in-context upskilling.

On productivity, the framework provides excellent diagnostic infrastructure but does not specify which tasks within the search marketing inventory must be redistributed, or where the boundaries between governance tiers should fall. The Work Engineering methodology implies a Maximalist work-as-product redesign, but the public framework does not provide the task-level specifications that would allow independent application. On pipeline, the bridging solutions are reasonable responses to the symptom but do not preserve the developmental work itself — they help organizations hire around the pipeline collapse rather than prevent it.

What Daugherty & Wilson Human + Machine would prescribe

Identify the missing middle: the zone of human-AI collaboration that lies between purely human work (creative strategy, client relationships) and purely machine work (bid optimization, programmatic placement). Within that middle zone, map role types from the 2018 framework. Trainers teach AI bid systems brand-specific patterns and client preferences. Explainers translate algorithmic creative recommendations to clients. Sustainers monitor AI-generated reporting and intervene when outputs drift. Apply the three fusion skills from the 2024 update: intelligent interrogation, judgment integration, and reciprocal apprenticing.

On productivity, the framework provides the best available vocabulary for what individual contributors do when they work well with AI. It does not distinguish governance intensity across tasks — all missing-middle work is treated as collaborative without specifying which collaboration requires what level of human control — and the role taxonomy is not designed to predict the productivity gain from redistributing tasks across tiers. On pipeline, the role taxonomy assumes the existence of analysts capable of training, explaining, and sustaining without theorizing how those analysts come to be.

The four prescriptions are answering different questions at different units of analysis: enterprise architecture (McKinsey), human-capital benchmarking with implied work redesign (Deloitte), individual-level role vocabulary (Daugherty & Wilson), and function-level governance redistribution (this paper). The empirical and theoretical case for the function-level prescription is set out in D5.