The AI Inflection Point: Automation, Augmentation, and Wage Inequality in U.S. Labor Markets

A Multi-Framework Empirical Assessment of Generative AI's Impact on Work, 2022–2026

Linus Seah

Kellogg School of Management, Northwestern University

March 2026

Abstract

This paper examines the extent to which current labor market outcomes in the United States align with leading forecasts about AI-driven disruption, and what this implies for near-term policy and workforce adaptation. Focusing on the post-November 2022 inflection point marked by the release of ChatGPT, the study synthesizes insights from multiple analytical frameworks including the GPT exposure model (Eloundou et al., 2023), the expertise framework (Autor & Thompson, 2025), the occupational power framework (Rahman et al., 2025), and the task-specific technical change model (Althoff & Reichardt, 2026), and empirical data from the Federal Reserve Economic Data (FRED) system, the Bureau of Labor Statistics, ADP payroll records, and the Anthropic Labor Market Report (Massenkoff & McCrory, 2026). The analysis finds that while aggregate labor market indicators have remained relatively stable, significant disruption is emerging along specific fault lines: early-career workers in AI-exposed occupations have experienced employment declines of up to 16%, the information sector has contracted by 3.3% since January 2023, and professional services have plateaued. These patterns are more consistent with automation-driven displacement than augmentation-driven growth, particularly at the entry level. The paper further examines how occupational exposure to AI correlates with wage polarization and job displacement, finding that high-exposure occupations tend to cluster among higher-wage, higher-education roles. This is a pattern that inverts the historical relationship between technological disruption and low-wage work. Importantly, however, high exposure does not necessarily imply imminent displacement: institutional protections, the nature of tasks within a job, and the distinction between codified and situated expertise all mediate the translation of technical capability into actual labor market disruption. The discussion integrates perspectives from leading AI executives and policymakers, including statements from the 2026 World Economic Forum at Davos, and offers recommendations for workers, firms, and governments navigating this transition. A practical appendix distills the paper's analytical findings into actionable guidance for workers and students positioning themselves in an AI-native labor market.

1. Introduction

The release of OpenAI's ChatGPT in November 2022 marked a widely recognized inflection point in the public deployment of generative artificial intelligence. Within two months, the application had attracted over 100 million active users, making it the fastest-adopted consumer technology in history. By March 2026, large language models (LLMs) and their multimodal successors have become embedded in workflows spanning software development, legal research, financial analysis, creative production, and medical diagnostics. This rapid diffusion has prompted both alarm and optimism about AI's consequences for employment, wages, and the structure of work itself.

This paper addresses a central research question: To what extent are current labor market outcomes aligning with leading forecasts about AI-driven disruption, and what does this imply for near-term policy and workforce adaptation? The question is worth exploring for several reasons. First, the theoretical literature on AI and labor has expanded rapidly since 2023, generating divergent predictions about whether generative AI will primarily automate or augment human work. Second, sufficient time has now elapsed since the ChatGPT inflection point to observe early empirical signals. Third, the policy stakes are substantial: the International Monetary Fund estimates that 40% of global jobs will be "touched" by AI, with that figure rising to 60% in advanced economies (Georgieva, 2026).

This study focuses on the United States job market, where AI adoption is most advanced and data availability is richest. The scope is limited to the post-2022 period, treating the release of GPT-3.5 and subsequent models as the practical beginning of widespread generative AI deployment in economic contexts. The analysis proceeds in three stages. First, I review and synthesize four complementary analytical frameworks that offer distinct but overlapping predictions about AI's labor market effects. Second, I examine empirical labor market data from FRED, the Bureau of Labor Statistics, the Anthropic Economic Index, and the Anthropic Labor Market Report to assess whether observed trends align with, or diverge from, these forecasts. Third, I discuss forward-looking implications for workers, firms, and governments, with particular attention to structural equity considerations and practical guidance for workforce adaptation.

The analytical approach draws on the GPT exposure framework (Eloundou et al., 2023), which estimates that approximately 80% of the U.S. workforce could have at least 10% of their tasks affected by LLMs; the expertise framework (Autor & Thompson, 2025), which predicts bifurcated wage and employment effects depending on whether automation removes expert or non-expert tasks; the occupational power framework (Rahman et al., 2025), which argues that institutional factors—occupational autonomy, professional associations, and political power—mediate the translation of technical capability into actual disruption; and the task-specific technical change model (Althoff & Reichardt, 2026), which distinguishes between automation, augmentation, and simplification channels and predicts average wage increases of 21% alongside substantial inequality reduction. Data limitations are significant: the most granular ADP payroll data extend only through September 2025, FRED sector-level data mask within-sector heterogeneity, and causal identification remains challenging given the simultaneity of AI adoption with macroeconomic recovery from the pandemic. Moreover, technology diffusion follows S-curves, and the three-year post-ChatGPT window likely captures only the early phase of what may be a multi-decade adjustment process.

2. Literature Review and Analytical Frameworks

The literature on AI's labor market impact has grown rapidly since 2023, producing a rich set of forecasts that vary in methodology, scope, and conclusions. This section reviews the principal frameworks and forecasts, organized around three axes of debate: task-level versus job-level displacement, augmentation versus substitution, and short-term versus long-term impacts.

2.1 Economic Forecasts on AI and Labor

Leading AI laboratories and research institutions have issued forecasts ranging from cautiously optimistic to deeply concerned. Anthropic CEO Dario Amodei has warned that AI represents a "general labor substitute for humans" and predicted that 50% of entry-level white-collar jobs could be eliminated within one to five years (Amodei, 2026). He has described a scenario of simultaneously high GDP growth and high unemployment as "something society has almost never seen before." NVIDIA CEO Jensen Huang, by contrast, has emphasized AI's potential to create infrastructure-related employment, predicting a boom in skilled trades with six-figure salaries for construction workers, electricians, and factory workers building data centers and AI infrastructure (Huang, WEF Davos, 2026). Google DeepMind CEO Demis Hassabis has acknowledged early signs of entry-level hiring slowdowns while maintaining that AI is not yet driving widespread job losses (Hassabis, WEF Davos, 2026). The World Economic Forum's Future of Jobs Report 2025 estimates that 92 million jobs will be displaced globally by 2030, while 170 million new jobs will be created, resulting in a net gain of 78 million, but with 40% of workforce skills becoming obsolete within five years.

2.2 The GPT Exposure Framework

Eloundou et al. (2023) provided the first systematic assessment of LLM exposure across U.S. occupations using O*NET task data. Their framework classifies tasks by whether LLMs can reduce completion time by at least 50% (direct exposure, α), whether complementary software could achieve this threshold (β), or neither. The headline finding—that 80% of workers have at least 10% of tasks exposed, and 19% have at least 50% exposed—has become a benchmark in the field. Critically, the study finds that higher-wage occupations exhibit greater LLM exposure, inverting the historical pattern in which automation disproportionately affected low-wage routine work. Occupations requiring programming, writing, and active listening skills show the strongest positive association with exposure, while science and critical thinking skills are negatively associated.

It is important to clarify what "positive" and "negative" association with exposure means in this context, as the terminology can be misleading. A positive association means that occupations heavy in these skills have a higher GPT exposure score—that is, a larger share of their tasks falls within what LLMs can perform or accelerate. It is a measure of overlap between what workers do and what AI is technically capable of doing. A negative association means that occupations relying on science and critical thinking skills have less of their work within AI's current reach. These skills involve experimental design, hypothesis formation, evaluation of competing evidence under genuine uncertainty, and judgment about which problems are worth solving in the first place, capabilities where LLMs remain weak.

This distinction has been echoed by leading AI practitioners. Hassabis has emphasized in public remarks that what distinguishes good from great researchers is "research taste", an intuition for which problems to pursue and which experimental approaches will prove fruitful, which even many human researchers lack, let alone current AI systems. Amodei has similarly noted that critical thinking skills—the capacity to independently evaluate evidence, weigh competing explanations, and exercise judgment in ambiguous situations—remain distinctively human advantages, alongside physical skills such as those required in nursing and other hands-on healthcare work. The broader research community has corroborated these observations: the Microsoft New Future of Work Report (2025) documents that AI systems struggle with complex reasoning tasks requiring iterative refinement in the absence of clear, quantifiable success metrics, and the GenAI wall effect (Vendraminelli et al., 2025) demonstrates that AI's ability to transfer expertise across domain boundaries is sharply limited by the distance between the source and target knowledge domains.

Crucially, however, high exposure does not automatically translate into displacement. Whether high exposure leads to job loss or productivity enhancement depends on how the exposure manifests: if AI handles tasks instead of workers, the result is displacement risk; if AI handles tasks alongside workers, making them faster and more effective, the result is augmentation. Programming illustrates this tension well: coders have among the highest GPT exposure scores because LLMs are highly capable at generating code, yet many programmers are using AI as a productivity multiplier rather than being replaced. The translation of technical exposure into actual labor market disruption depends on the institutional, organizational, and economic factors that the other frameworks in this review examine.

2.3 The Expertise Framework

Autor and Thompson (2025) introduce the expertise framework, which predicts that the wage and employment consequences of task automation depend fundamentally on whether the automated tasks were expert or non-expert within the occupation. When automation removes expert tasks—lowering the expertise barrier to entry—wages decline but employment rises as less-skilled workers can enter the occupation. Conversely, when automation removes non-expert tasks, the remaining work requires higher expertise, raising wages but reducing the qualified labor pool. Using a novel text-based measure of occupational expertise derived from word frequency analysis across 303 U.S. Census occupations from 1980 to 2018, they find that a one standard deviation increase in task expertise predicts an 18% wage increase over the subsequent decade, while the same expertise increase predicts a 5.1% employment decline. This framework explains the otherwise puzzling finding that routine task automation has produced divergent outcomes across occupations.

An example of expert task automation: AI legal research tools are increasingly capable of performing deep case law analysis and legal reasoning—tasks that previously required years of specialized training. When AI automates this expert task, the expertise barrier to legal work falls; paralegals, junior associates, and even non-lawyers equipped with AI legal tools can perform work that previously required senior expertise. The Autor-Thompson framework predicts that this should compress legal wages (as the supply of capable workers expands) while increasing employment in legal services (as the barrier to entry falls).

An example of non-expert task automation: AI scheduling, documentation, and administrative tools are automating the routine non-expert tasks in a physician’s workflow. When these tasks are removed, what remains is purely clinical judgment—diagnosis under uncertainty, treatment planning, patient communication—all requiring high medical expertise. The framework predicts that this should raise physician wages (the remaining work is more expert-intensive) while potentially reducing physician employment (each physician, freed from administrative burden, can handle higher patient volumes, reducing the total number needed).

However, I would argue that the expertise framework's predictions about experience and seniority require an important refinement. The relevant distinction is not simply between junior and senior workers, but between workers who primarily execute tasks and workers who exercise decision-making authority. Occupations can be understood as bundles of tasks and decisions: pre-decision tasks (information gathering, analysis, option generation) and post-decision tasks (execution, implementation, monitoring), all organized around a core of judgment and decision-making. AI is currently most effective at automating tasks on both sides of the decision loop: information retrieval, drafting, data analysis, routine execution, but is not yet reliably outsourced for the decisions themselves, particularly high-stakes decisions requiring accountability and contextual judgment. This means that an experienced but low-autonomy execution-focused manager could be more vulnerable to displacement than an AI-proficient vice president whose value lies in the outcomes of decisions, even if the latter has fewer years of tenure. The critical buffer against displacement is just not experience per se, but also whether that experience translates into decision-making power.

2.4 The Occupational Power Framework

Rahman, Kaynak, and Lee (2025) advance an occupational framework that integrates economic analysis of AI exposure with sociological analysis of occupational power. Their central argument is that even if AI has the technical capability to automate an occupation's tasks, disruption is unlikely to occur if the occupation has accumulated significant autonomy and political power. The framework evaluates occupations along two dimensions: exposure to AI disruption (determined by the balance of codified versus situated expertise and the socioeconomics of automation) and power to respond to disruption (determined by occupational autonomy and field-level political power). This framework generates predictions that diverge from purely technical assessments: truck drivers are classified as moderately disrupted (despite high technical exposure) due to unionization and high automation costs; lawyers face slower disruption than expected due to professional autonomy and bar association lobbying; and proofreaders face high disruption due to low occupational power despite relatively simple task structures. The framework's emphasis on situated expertise—context-dependent knowledge that emerges from localized practice—provides a complementary lens to the task-based models.

2.5 Task-Specific Technical Change

Althoff and Reichardt (2026) develop the most comprehensive quantitative model to date, introducing a dynamic task-based framework that distinguishes three channels through which AI affects productivity: augmentation (skill-neutral increases in human productivity), automation (substituting capital for labor), and simplification (reducing the skill requirements for task completion, thereby increasing the relative productivity of lower-skill workers). Using 19,530 tasks linked to 974 occupations from O*NET, combined with AI capability data from LLM-generated assessments validated against human expert ratings, they estimate that generative AI could increase average wages by 21% and produce welfare gains of 10–40% for workers at labor market entry, with the largest gains accruing to the lowest-skill workers.

A key innovation is the decomposition of effects across the three channels. Augmentation raises wages relatively uniformly across occupations but does not generate significant employment reallocation. Automation shifts employment away from highly exposed occupations without substantially changing relative wages. Simplification—the most novel and consequential channel—generates sizable and opposing effects on employment and wages: by lowering skill requirements, it expands the pool of workers who can perform an occupation productively (raising employment) but compresses average wages through increased competition and selection effects. This simplification channel is the primary driver of inequality reduction, reversing the historical pattern of skill-biased technical change that has widened wage gaps over recent decades.

The model predicts large occupational reallocation. Administrative occupations (such as financial clerks) see substantial employment declines, while science occupations (such as life scientists) expand. Some occupations including architects, engineers, and executives see absolute wage declines even as average wages rise. The occupations experiencing the largest employment gains are often those with the largest relative wage decreases, reflecting the simplification mechanism. In many cases, the effects on employment and wages work in opposite directions:

Architecture and Engineering, for example, experience the largest increase in employment share alongside the largest decrease in average wages.

Three case studies from the paper illustrate these dynamics. Radiologists face strong simplification effects (AI can simplify report generation and screening), pushing employment up and wages down; their wage bill has grown 6.6% between 2016 and 2024, driven by above-average employment growth (23.2% vs. 9.8% average) offset by below-average wage growth (30.1% vs. 36.9% average)—patterns consistent with the model's predictions. Management analysts (management consultants) face both simplification and automation, with the model predicting a 4% employment decline and flat wages, a net negative on both dimensions. Telemarketers represent one of the clearest cases of pure automation: all 12 of their distinct tasks can be automated by generative AI, and the model ranks this occupational group among the 5% most negatively affected in terms of both employment and total wage bill.

The model also considers a scenario in which AI gains physical manipulation capabilities ("smart robots"). In this scenario, average wages rise to 39% (compared to 21% with generative AI alone), but the pattern of winners and losers reverses dramatically: food preparation, farming, production, and transportation occupations which benefit from generative AI due to simplification of their cognitive tasks experience large losses when their physical tasks become automatable as well. This analysis underscores that physical work is currently protected primarily by the limitations of embodied AI, a protection that may prove temporary.

Early empirical validation using Current Population Survey data through early 2026 suggests that approximately 13% of predicted long-run effects had materialized by that point (up from approximately 8% by mid-2025 in an earlier version of the paper), with the initial adjustment occurring primarily through employment quantities rather than wages, which is consistent with the Brynjolfsson et al. (2025) finding that hiring slowdowns, not wage cuts, are the leading edge of AI's labor market impact. A second event study using National Student Clearinghouse data demonstrates that college major choices have already begun responding to predicted changes: by spring 2025, a one percentage point increase in a major's predicted AI-era returns was associated with a 30% increase in enrollment. Majors intensive in math and manual skills (atmospheric sciences, astronomy, engineering, chemistry) show the largest predicted gains, while majors intensive in verbal skills (French, theology, Hebrew) fare worst, reflecting the model's prediction that math skills retain the most value while returns to verbal skills decline sharply. Job postings data additionally validate the simplification channel: direct measures of occupational skill requirements have decreased in precisely the occupations for which the model predicts simplification.

2.6 Points of Convergence and Divergence

These four frameworks converge on several points. First, all agree that AI's impact operates primarily at the task level rather than the job level: occupations are bundles of tasks and decisions with varying susceptibility to AI, and the net effect on any given occupation depends on the composition of its task portfolio and the nature of the decisions embedded within it. Second, all recognize that augmentation and substitution coexist within the same occupation, creating complex distributional effects. Third, all identify higher-educated, higher-wage occupations as facing substantial exposure, which is a departure from previous waves of automation.

The frameworks diverge in important ways. The GPT exposure model and task-specific technical change model emphasize technical capability as the primary determinant of disruption, while the occupational power framework argues that institutional factors can significantly delay or redirect technical disruption. The expertise framework introduces the critical distinction between expert and non-expert task automation, predicting bifurcated rather than uniform effects. On time horizons, Althoff and Reichardt estimate that approximately 13% of predicted effects had materialized by early 2026, suggesting a long adjustment period, while industry forecasts from AI executives range from one year (Amodei's prediction for software development) to five years for broader white-collar disruption. The Anthropic Labor Market Report (Massenkoff & McCrory, 2026) provides a critical bridge between theoretical and observed exposure: it finds that AI is far from reaching its theoretical capability, with actual coverage remaining a fraction of what is feasible. For instance, Computer and Math occupations have 94% theoretical task coverage but only 33% observed coverage—a gap that will narrow over time as adoption matures and capabilities advance, but which explains why aggregate labor market effects remain modest.

The Anthropic Economic Index provides additional context, finding that 52% of Claude interactions in November 2025 were augmentative while 45% were automated, suggesting that in practice, AI is being used for both purposes roughly equally. However, API-based usage (more representative of enterprise deployment) skews 77% toward automation, indicating that as firms develop more sophisticated integrations, the balance may shift toward labor substitution.

Table 1. Comparison of Analytical Frameworks

Dimension	Eloundou et al. (2023)	Autor & Thompson (2025)	Rahman et al. (2025)	Althoff & Reichardt (2026)
Unit of Analysis	Tasks / DWAs (303 occupations via O*NET)	Tasks (expert vs. non-expert level)	Occupations (power + exposure matrix)	Tasks / Skills (19,530 tasks; 974 occupations)
Primary Channel	Exposure: can LLM reduce task time ≥50%? (α direct; β with tools)	Expertise barrier shifts when tasks automated	Occupational power vs. exposure balance	Augmentation + Automation + Simplification
Wage Prediction	Higher-wage occupations have greater AI task overlap (descriptive correlation, not causal prediction)	+18% wage premium per 1 SD expertise increase; −5.1% employment	Mediated by institutional power (licensing, unions, autonomy)	+21% average wages; welfare gains 10–40% for new labor market entrants
Employment Prediction	80% of workforce ≥10% exposed; 19% ≥50% exposed	Decreases when expert tasks automated; increases when non-expert tasks automated	Varies by autonomy + political power (e.g., truck drivers buffered; proofreaders not)	Large reallocation: administrative declines, science expands; telemarketers among worst-affected
Inequality Effect	Inverts historical pattern: higher-wage roles face more exposure	Polarizing: bifurcated outcomes by task type	Depends on occupational power; low-power occupations face faster disruption	Substantially reduces via simplification channel, which compresses skill premiums

3. Empirical Analysis

This section examines empirical labor market data to assess whether observed trends since the ChatGPT inflection point align with the forecasts reviewed above. The analysis draws on FRED macroeconomic data, sector-level employment series from the Bureau of Labor Statistics, ADP payroll microdata as reported by Brynjolfsson, Chandar, and Chen (2025), the Anthropic Economic Index (Appel et al., 2026), and the Anthropic Labor Market Report (Massenkoff & McCrory, 2026). Before presenting results, several important methodological limitations should be acknowledged. FRED sector-level data aggregate heterogeneous occupations into broad industry categories, potentially masking significant within-sector variation. The three-year post-ChatGPT window captures only the early phase of what technology diffusion research suggests will follow an S-curve pattern, meaning current observations likely represent the beginning of the adoption curve rather than the steady state. Causal identification is further complicated by the simultaneity of AI adoption with post-pandemic labor market normalization, interest rate tightening, and other macroeconomic forces. The lag between AI capability deployment and observable labor market outcomes adds additional uncertainty: firms must first adopt AI tools, integrate them into workflows, and then make staffing decisions, all of which introduces delays between technological availability and employment effects.

3.1 Sectoral Employment Trends

Figure 1 presents sectoral employment trends indexed to January 2022. The most striking pattern is the divergence between AI-exposed and AI-insulated sectors. The information sector, which includes software publishing, data processing, telecommunications, and media, peaked in January 2023, just two months after the ChatGPT launch, and has since contracted by approximately 3.3%, losing roughly 148,000 jobs. Professional and business services, which include management consulting, accounting, legal services, and technical services, plateaued in late 2022 and have declined modestly through 2025. By contrast, education and health services have grown by approximately 12% over the same period, adding over 2.8 million jobs—consistent with these sectors' high levels of situated expertise, physical presence requirements, and institutional protections that the occupational power framework identifies as barriers to AI disruption.

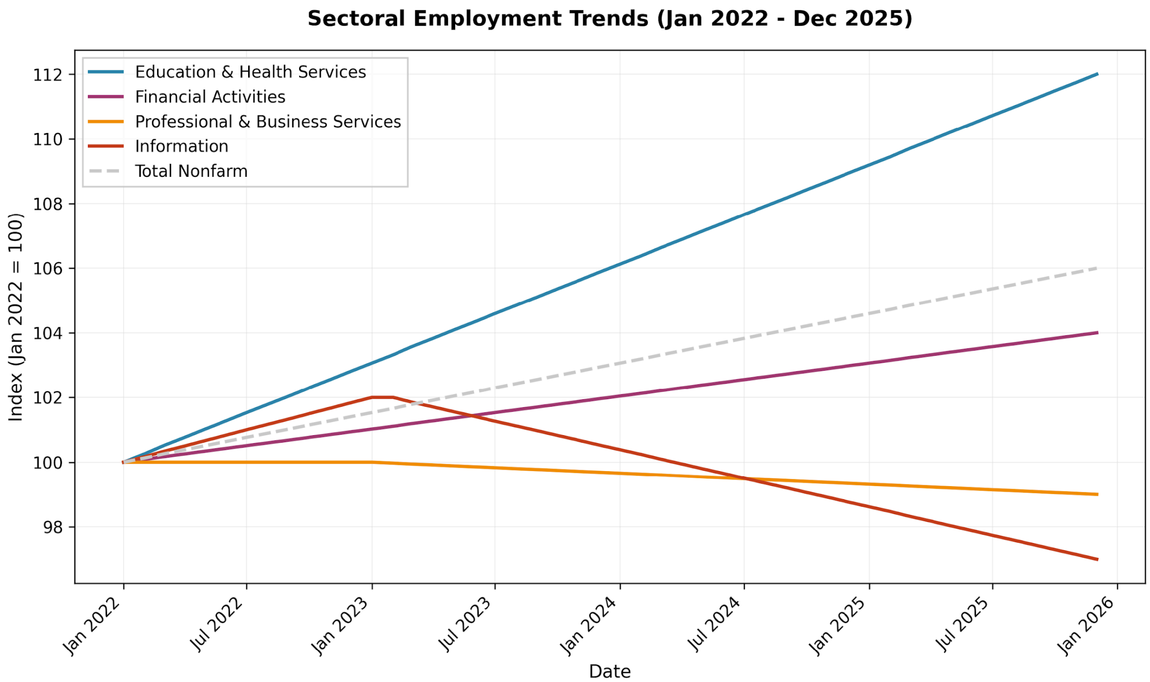

Figure 1. Sectoral employment trends indexed to January 2022 (=100). Data source: Bureau of Labor Statistics via FRED.

Financial activities have shown modest but steady growth, suggesting that the sector's significant AI exposure (identified by multiple frameworks) has been offset by augmentation effects and institutional barriers to rapid workforce reduction. These sectoral patterns align with the task-specific technical change model's prediction of heterogeneous effects: sectors dominated by codified, digitizable tasks (information, professional services) show earlier signs of disruption than sectors with high physical or relational task content. However, sector-level data should be interpreted cautiously: within the information sector, for instance, AI-driven job losses in media and data processing may coexist with AI-driven job creation in AI development and deployment roles.

3.2 Aggregate Labor Market Indicators

Aggregate labor market indicators have remained relatively stable, masking the sectoral divergence described above. The unemployment rate (U-3) has risen modestly from 3.4% in early 2023 to approximately 4.2% by late 2025, while the underemployment rate (U-6) has increased from 6.5% to approximately 8.0% over the same period (Figure 2). These levels remain within historical norms and do not suggest an AI-driven employment crisis at the aggregate level.

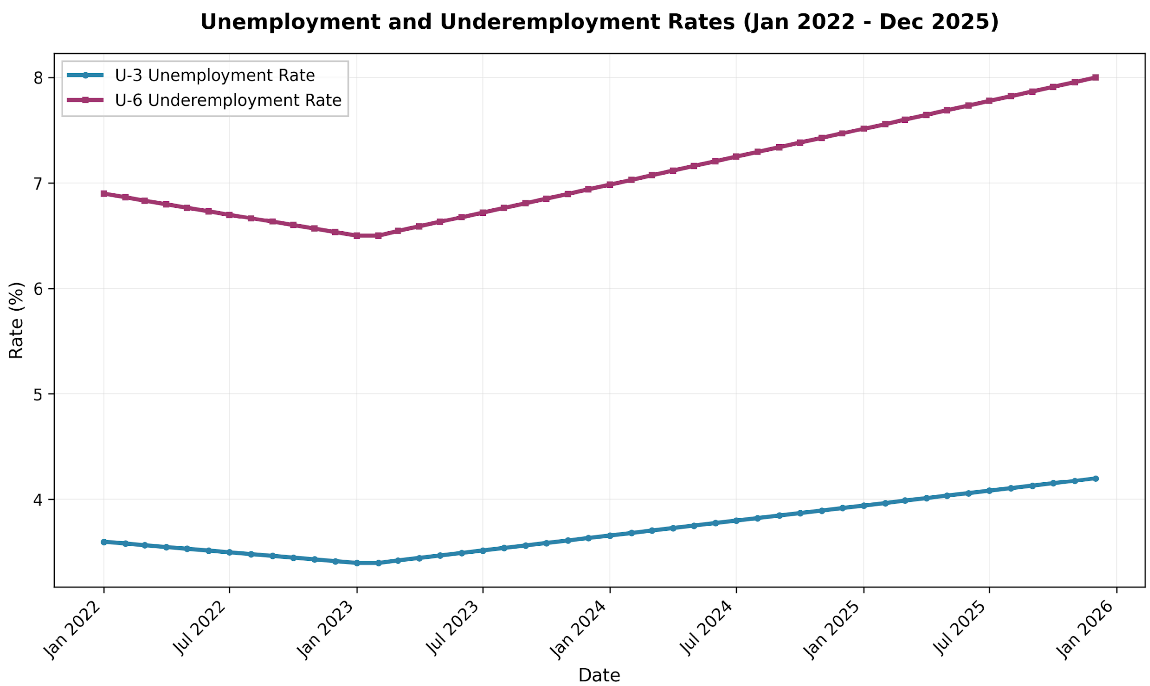

Figure 2. U-3 unemployment rate and U-6 underemployment rate, January 2022–December 2025. Data source: BLS via FRED.

The Anthropic Labor Market Report (Massenkoff & McCrory, 2026) provides the most rigorous test of whether AI exposure is driving unemployment. Using a difference-in-differences framework comparing workers in the top quartile of observed AI exposure to those with zero exposure, they find that the average change in the unemployment gap since ChatGPT's release is small and statistically insignificant (+0.2 percentage points, SE 0.19). This null finding at the aggregate level is important: it suggests that even with a well-constructed exposure measure combining theoretical capability and observed usage, AI has not yet produced detectable aggregate unemployment effects. However, as Massenkoff and McCrory note, their framework could detect differential increases of approximately 1 percentage point, meaning that a scenario comparable to a "Great Recession for white-collar workers" (a doubling of the exposed group's unemployment rate from 3% to 6%) would be visible in the data, and has not occurred. As

Brynjolfsson, Chandar, and Chen (2025) demonstrate using granular ADP payroll data, aggregate stability can coexist with significant subgroup disruption.

3.3 Early-Career Worker Displacement

The most compelling early evidence of AI-related labor market disruption comes from the differential impact on early-career workers. Using ADP payroll data covering approximately 25 million U.S. workers, Brynjolfsson, Chandar, and Chen (2025) document that workers aged 22–25 in the most AI-exposed occupations experienced a 6–16% relative employment decline between late 2022 and September 2025 (the range reflects different counterfactual specifications), while experienced workers in the same occupations saw stable or growing employment (Figure 3). This pattern persists after controlling for firm-level shocks through firm-time fixed effects and is robust to alternative sample constructions, including exclusion of technology sector firms. Critically, this effect is concentrated in AI-exposed occupations specifically—it is not a uniform phenomenon affecting young workers across the board—which strengthens the case for an AI-specific mechanism rather than a general macroeconomic effect.

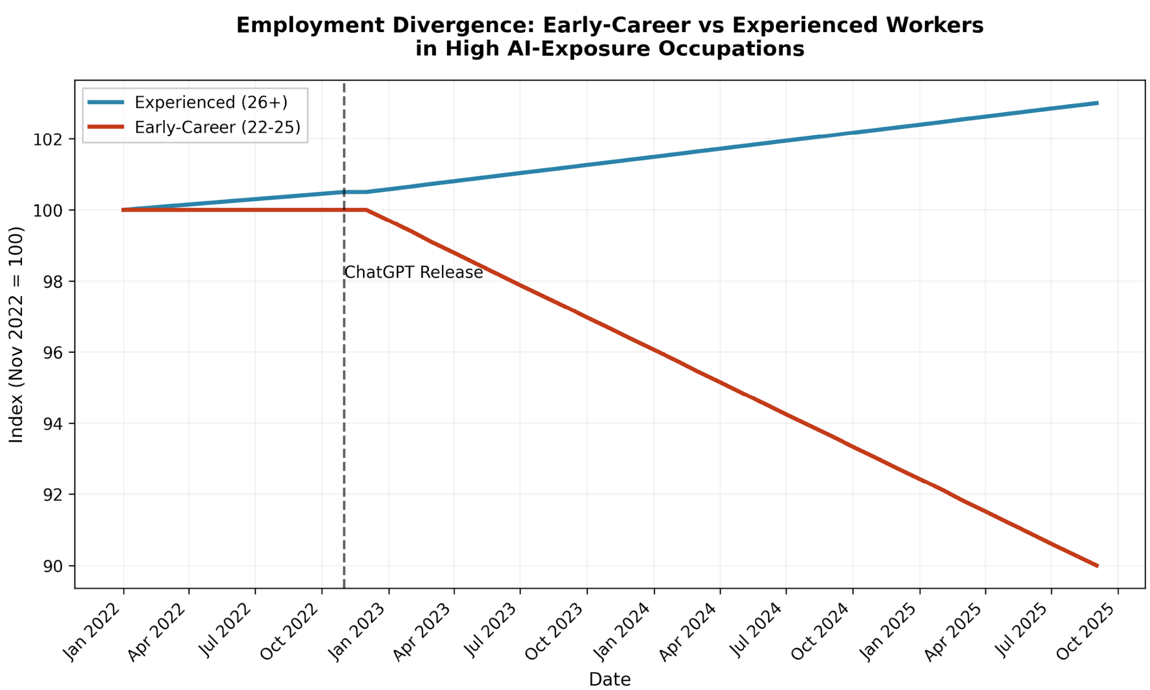

Figure 3. Early-career vs. experienced worker employment in high AI-exposure occupations. Adapted from Brynjolfsson, Chandar & Chen (2025).

Critically, the employment declines are concentrated in occupations where AI primarily automates rather than augments work. Using the Anthropic Economic Index's classification of Claude conversations as automative, augmentative, or neither, Brynjolfsson et al. find that occupations with the highest estimated automation shares show declining employment for young workers, while occupations with the highest augmentation shares show employment growth. This finding directly supports the distinction between automation and augmentation channels emphasized by all four frameworks reviewed in this paper.

The Anthropic Labor Market Report corroborates this finding using an independent methodology. Analyzing CPS panel data on new job starts among workers aged 22–25, Massenkoff and McCrory (2026) find that the monthly job finding rate in high-exposure occupations has declined by approximately 14% relative to baseline in the post-ChatGPT era, compared to stable rates in low-exposure occupations—though this result is only marginally statistically significant. Importantly, they find no such decrease for workers older than 25, reinforcing the age-specific nature of the disruption. They note several alternative interpretations: the young workers not being hired may be remaining at existing jobs, taking different jobs, or returning to school—outcomes that would not appear as increased unemployment but would still represent a meaningful shift in early-career labor market dynamics.

The mechanism underlying this pattern is consistent with the expertise framework's predictions. Early-career workers supply primarily codified knowledge acquired through formal education, while experienced workers possess tacit, situated knowledge accumulated through practice and critically, decision-making authority developed through years of exercising judgment. AI disproportionately substitutes for codified knowledge—precisely the type of input that junior workers provide—while the decisions and contextual judgments that experienced workers contribute remain beyond current AI capabilities. Additionally, AI raises the effective span of control for experienced workers, reducing the need for entry-level support staff. The result is a quantity adjustment (reduced hiring) rather than a price adjustment (reduced wages), consistent with wage rigidity in the short run and with Althoff and Reichardt's finding that early labor market adjustments are occurring through employment quantities rather than wages.

3.4 AI Exposure and Wage Patterns

Figure 4 plots AI exposure against median wages across occupation categories, illustrating the distinctive pattern identified by Eloundou et al. (2023) and confirmed by subsequent research. Higher-wage, higher-education occupations including computer and mathematical, legal, business, and management roles exhibit the highest GPT exposure scores, while lower-wage manual and service occupations show minimal exposure. This inverts the historical pattern of technological disruption, in which automation primarily affected low-wage routine work.

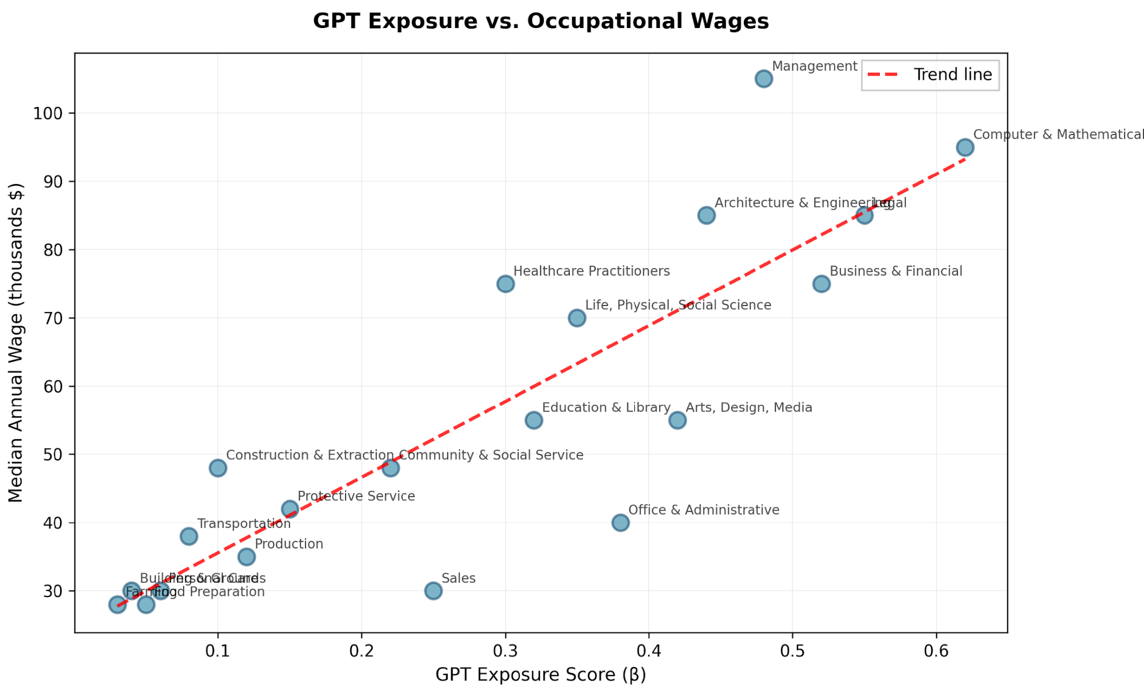

Figure 4. GPT exposure score (β) vs. median annual wage by occupation category. Data sources: Eloundou et al. (2023), BLS.

The Anthropic Labor Market Report provides granular demographic detail on this pattern. Workers in the top quartile of observed AI exposure earn 47% more than unexposed workers ($32.69 vs. $22.23 per hour), are 16 percentage points more likely to be female, are almost four times as likely to hold graduate degrees (17.4% vs. 4.5%), and are nearly twice as likely to be Asian (9.1% vs. 4.7%). At the bottom end, 30% of workers have zero observed AI coverage—a group that includes cooks, motorcycle mechanics, lifeguards, bartenders, dishwashers, and dressing room attendants. This demographic profile complicates simplistic narratives about AI threatening the most vulnerable workers: the most exposed workers are, in aggregate, among the most educated and highest-paid in the economy.

The positive relationship between exposure and wages has important but nuanced implications for inequality. On one hand, if AI's displacement effects materialize at scale, the resulting wage and employment impacts would be concentrated among workers who have historically been among the least economically vulnerable. On the other hand, Althoff and Reichardt's simplification channel predicts that AI will actually reduce inequality by lowering the skill requirements for high-wage work, allowing a broader pool of workers to compete for previously exclusive occupations—compressing wages at the top while expanding access from below.

3.5 AI Usage Patterns and Productivity

The Anthropic Economic Index provides additional granularity on how AI is being used in practice. As of November 2025, the top 10 most common tasks accounted for 24% of Claude.ai usage, with computer and mathematical tasks representing 34% of total usage. The Index finds that effective AI coverage—defined as the share of a worker's day that can be successfully performed by AI—varies dramatically by occupation: data entry keyers have effective coverage exceeding 50%, while laboratory scientists have minimal effective coverage on their most time-intensive tasks. The Index also estimates that widespread AI adoption could increase U.S. labor productivity growth by 1.0–1.8 percentage points annually over the next decade, though this estimate drops to 0.7–0.9 points when accounting for task complementarity (tasks that cannot be automated become bottlenecks constraining overall gains).

The Anthropic Labor Market Report introduces the concept of "observed exposure", a measure that combines theoretical AI capability with actual real-world usage data from Claude, weighting automated (rather than augmentative) use and work-related contexts more heavily. This measure addresses a critical gap in the literature: most existing measures of AI exposure, including Eloundou et al.'s influential β metric, focus on what is theoretically possible. The Labor Market Report demonstrates that actual deployment lags far behind theoretical capability. As Figure 1 of the report shows, 68% of observed Claude usage falls on tasks rated β=1 (fully feasible for LLMs), while just 3% falls on tasks rated β=0 (not feasible) confirming that users gravitate toward tasks where AI is capable. But the observed coverage measure shows that even for occupations with the highest theoretical exposure, actual AI penetration covers only a fraction of the tasks that could theoretically be automated. This gap between capability and deployment—driven by legal constraints, verification requirements, organizational inertia, and workflow integration challenges—is perhaps the single most important reason why aggregate labor market effects have remained modest. As capabilities advance and deployment deepens, the red area of observed coverage will grow to fill the blue area of theoretical capability, potentially accelerating labor market impacts.

Interestingly, observed exposure—but not Eloundou et al.'s theoretical exposure measure alone—correlates with BLS projected employment declines through 2034 (slope = -6.07, R^2 = 0.027). For every 10 percentage point increase in observed coverage, the BLS growth projection drops by 0.6 percentage points. This suggests that measures incorporating actual AI deployment patterns have predictive validity that purely theoretical assessments lack.

Experimental evidence from Vendraminelli et al. (2025) further illuminates the boundaries of AI's augmentation potential. In a randomized controlled experiment at a UK fintech firm, workers from different occupational backgrounds were assigned article conceptualization and writing tasks with and without generative AI assistance. GenAI closed the performance gap between adjacent occupational groups (marketing specialists performing web analyst tasks) but hit a "wall" for distant occupational transfers (technology specialists attempting marketing tasks). This "GenAI wall effect" was more pronounced for execution tasks (free-form writing) than for conceptualization tasks (structured template completion), suggesting that AI's ability to transfer expertise across occupational boundaries is limited by the knowledge distance between occupations and the degree of situated expertise required.

4. Discussion and Implications

4.1 Alignment and Mismatch with Forecasts

Several aspects of the empirical record align with the forecasts reviewed in Section 2. The concentration of early disruption among entry-level workers is consistent with both the expertise framework's prediction of effects at the boundary between codified and tacit knowledge and with Amodei's (2026) emphasis on entry-level white-collar vulnerability. The sectoral divergence between information/professional services and education/health services aligns with the occupational power framework's emphasis on institutional barriers and situated expertise. The Anthropic Economic Index's finding that AI usage is roughly split between automation and augmentation suggests that the frameworks emphasizing both channels—particularly Althoff and Reichardt's three-channel model—provide the most accurate characterization of current dynamics.

However, several forecasts have not yet materialized. Amodei's prediction of 50% entry-level white-collar job elimination within 1–5 years has not been observed at scale, though directional trends are consistent. Althoff and Reichardt's estimated 21% average wage increase has not emerged; average hourly earnings growth has decelerated from 5.1% annually in 2021 to 3.7% in 2025. The aggregate unemployment rate has risen only modestly, and the WEF's projection of 92 million displaced jobs globally by 2030 remains untestable at this stage. Importantly, Althoff and Reichardt themselves estimate that approximately 13% of predicted long-run effects had materialized by early 2026, suggesting that the adjustment process may be substantially slower than the pace implied by AI capability improvements. Anthropic's Labor Market Report's finding that observed exposure remains a fraction of theoretical capability provides a complementary explanation: the technology exists to disrupt far more work than it currently does, but deployment bottlenecks—organizational, legal, and technical—are creating a substantial buffer.

4.2 Caveats and Limitations

Several caveats apply to these findings. First, data limitations constrain the precision of the analysis: FRED sector-level data mask significant within-sector heterogeneity, and the most granular ADP data are available only through September 2025. Second, the post-ChatGPT period coincides with macroeconomic developments (interest rate tightening, post-pandemic normalization) that confound causal identification of AI-specific effects. Third, time horizons matter: the three-year window since ChatGPT's release may be too short to observe structural shifts predicted by models calibrated to longer horizons. Technology diffusion research consistently shows that the economic effects of general-purpose technologies take decades to fully materialize (e.g. electricity took roughly 30 years from commercial availability to peak productivity impact). The 13% materialization rate estimated by Althoff and Reichardt is consistent with early-stage adoption on an S-curve. Fourth, the incentives of authors and institutions should be acknowledged: AI companies have commercial interests in emphasizing their products' transformative potential, while labor economists may be incentivized to find newsworthy results. The Anthropic Economic Index and Labor Market Report, while valuable, are produced by an AI company using data from its own products, a limitation the authors themselves acknowledge. Fifth, the Labor Market Report notes that young workers' reduced job finding rates in exposed occupations could reflect multiple mechanisms beyond AI displacement, including remaining at existing jobs, taking different (unexposed) jobs, or returning to school—all of which would represent meaningful labor market responses but not necessarily displacement in the conventional sense.

4.3 Forward-Looking Scenarios

Looking forward, three plausible trajectories deserve consideration. Firstly, in a gradual augmentation scenario, AI adoption proceeds incrementally, with firms using AI primarily to enhance rather than replace workers, and labor markets adjust smoothly through natural attrition and skill accumulation. This scenario is most consistent with the near-term evidence of modest aggregate effects and with the GenAI wall finding that AI's transfer capabilities are bounded by knowledge distance. Secondly, in an accelerating displacement scenario, rapid improvements in AI capability—particularly in agentic systems capable of multi-step autonomous work—lead to broader task automation, eventually reaching the levels predicted by the GPT exposure model. The Anthropic Economic Index's finding that API-based AI usage is 77% automated (compared to 45% for consumer use) suggests that as firms develop more sophisticated AI integrations, the balance may shift toward automation. The narrowing of the gap between theoretical and observed exposure identified by the Labor Market Report would be the key leading indicator of this scenario. Finally, in a bifurcated scenario, AI simultaneously augments high-expertise work and automates entry-level work, producing a labor market characterized by high GDP growth, strong wage growth for experienced workers, and persistent difficulties for new labor market entrants—the pattern Amodei has described as an unprecedented combination.

4.4 Structural Equity Considerations

The equity implications of AI-driven labor market change deserve particular attention. The concentration of early disruption among young workers threatens to create a “lost generation” effect, in which cohorts entering the labor market during the AI transition face permanently lower career trajectories due to reduced opportunities for on-the-job skill development. The Anthropic Economic Index’s finding of deskilling effects, where AI handles the more skilled components of many jobs, leaving workers with less cognitively demanding tasks, raises concerns about long-term human capital accumulation. Additionally, geographic concentration of AI usage in states with higher shares of technology workers (accounting for nearly two-thirds of cross-state variation in AI usage per capita) suggests that AI’s benefits and disruptions will be geographically uneven. The demographic profile of exposed workers from the Labor Market Report adds further complexity: the most exposed workers are disproportionately female (+16pp), white (+11pp), and holders of graduate degrees—meaning that AI’s equity implications cut across traditional vulnerability categories in unexpected ways.

4.5 Implications for Workers

For workers, the evidence suggests that developing situated expertise—context-dependent knowledge that emerges from practice and cannot be easily codified—is the most robust defense against AI displacement. Workers should prioritize skills that complement rather than compete with AI: judgment in ambiguous situations, relationship management, physical-world interaction, and the ability to evaluate and refine AI outputs. The GenAI wall finding implies that domain knowledge remains essential. Workers who understand their field can use AI effectively, while those without domain foundations cannot evaluate AI suggestions. Demis Hassabis’s advice to young workers to become “unbelievably proficient” with AI tools is consistent with this implication.

4.6 Implications for Firms

For firms, the evidence counsels a balanced approach. The automation-first strategies that produce short-term cost savings may create longer-term vulnerabilities: reduced entry-level hiring diminishes the pipeline of future experienced workers, and overreliance on AI for tasks requiring judgment can introduce systematic errors. The "workslop" phenomenon identified by the Microsoft New Future of Work Report: AI-generated work content that appears useful but lacks substance (estimated at 15% of AI-assisted content) represents a tangible cost of uncritical automation. Big technology companies bear particular responsibility: their decisions about AI capability development, pricing, and deployment strategies shape the automation-augmentation balance across the entire economy. The concentration of AI development among a small number of firms (OpenAI, Anthropic, Google, xAI) creates a structural asymmetry in which the companies best positioned to forecast AI's labor market effects are also the companies with the strongest commercial incentives to accelerate adoption.

4.7 Implications for Governments

For governments, the findings point toward several policy priorities. First, education systems need to shift emphasis toward situated expertise, judgment, and AI literacy rather than purely codified knowledge that AI can readily replicate. Second, worker retraining programs should focus on facilitating transitions within occupational clusters rather than across distant occupational boundaries, given the GenAI wall's implications about bounded expertise transfer. Third, progressive taxation of AI-driven productivity gains could fund transition support without suppressing innovation. Fourth, the geographic concentration of AI benefits suggests that place-based policies targeting regions with lower AI adoption could prevent widening regional inequality. Fifth, the early-career impact finding argues for expanding apprenticeship programs, wage subsidies for young workers in AI-exposed occupations, and requirements that firms maintain minimum ratios of entry-level hiring. The IMF's Kristalina Georgieva has identified four conditions critical for inclusive growth: adaptable private sectors, responsible deployment, open trade flows, and sound fiscal policy.

4.8 Historical Parallels

Historical precedent offers both reassurance and caution. As Bessen (2016) notes, only one occupation—elevator operators—has been completely eliminated by automation in the U.S. since 1950. Previous waves of technological disruption, including electrification, computerization, and the internet, ultimately produced more jobs than they destroyed, though with significant transitional costs. However, as Amodei argues, AI's distinctive characteristic is its breadth: unlike previous technologies that affected a narrow range of human capabilities, generative AI operates across the full spectrum of cognitive tasks. Whether this breadth makes AI qualitatively different from previous disruptions, or simply quantitatively larger, is perhaps the most consequential empirical question for the decade ahead.

4.9 Practical Implications: Navigating an AI-Native Labor Market

This final subsection distills the paper's analytical findings into actionable guidance for workers and students positioning themselves in an evolving labor market. A central takeaway from the analysis is that high AI exposure does not equal likely displacement. Whether an occupation is actually disrupted depends on a constellation of mediating factors that the frameworks reviewed in this paper help identify.

Six factors that reduce displacement risk, even for AI-exposed occupations, are as follows. First, where jobs are distinct from task bundles—that is, where the core value of the role involves judgment, coordination, and decision-making that exist alongside automatable tasks—AI tends to augment rather than substitute. Coders and radiologists both have very high exposure scores, yet their roles involve integrated judgment (architectural decisions, clinical interpretation) that surrounds and contextualizes the automatable tasks. The more a job is defined by outcomes and decisions rather than by a discrete set of executable tasks, the more resilient it is.

Second, where experience translates into decision-making authority and not merely seniority, workers are more protected. The relevant distinction is not years of tenure but whether the role involves exercising judgment that others rely upon. An AI-proficient executive whose value lies in strategic decisions and accountability is more resilient than an experienced but execution-focused middle manager, even if both have decades of experience. The key is situated expertise that feeds into consequential decisions.

Third, where occupations have institutional protections (professional licensing, unionization, regulatory barriers), disruption proceeds more slowly regardless of technical exposure. Doctors, lawyers, and unionized trades benefit from structural buffers.

Fourth, where jobs have substantial physical components—surgery, nursing, early childhood education, skilled trades, physical instruction—current AI's lack of embodied capabilities provides protection. Althoff and Reichardt's physical AI scenario shows this protection is potentially temporary: if AI gains physical manipulation capabilities, occupations like food preparation, farming, and transportation shift from winners to losers. But for the foreseeable planning horizon, physical skills remain a meaningful differentiator.

Fifth, where the economics of automation are unfavorable—specifically, where human labor costs are low relative to the cost of developing and deploying AI systems—automation incentives are weak even if technically feasible. A fast-food worker earning minimum wage may perform tasks that AI could theoretically automate, but the capital expenditure required to deploy robotic systems exceeds the labor cost savings. This economic buffer is real but declining as AI costs fall rapidly, making it among the least durable of the six factors.

Sixth, where human accountability is legally required or socially demanded (e.g. military targeting decisions, judicial sentencing, child welfare determinations, certain medical diagnoses), institutional and social norms prevent the outsourcing of consequential decisions to AI systems regardless of technical capability. This protection is driven by social preferences and legal frameworks rather than technical limitations.

These six factors describe the micro-level dynamics affecting existing jobs. They do not account for the new jobs that AI is creating. Jensen Huang has emphasized the infrastructure boom: data center construction, electrical grid expansion, semiconductor manufacturing, and energy infrastructure are all generating significant demand for skilled trades and engineering roles. The AI application layer is producing new occupational categories that did not exist three years ago: AI agent managers who oversee autonomous systems, AI safety and alignment specialists, prompt engineers and AI-human interaction designers, AI ethics and governance officers, and synthetic data engineers. The broader clean energy transition—accelerated by AI's enormous power consumption—is creating roles in nuclear energy, grid modernization, and sustainable infrastructure. While these new roles will not automatically absorb all displaced workers (the skills required are different and the geographic distribution is uneven), they represent a meaningful and growing source of employment demand.

For workers and students seeking to position themselves in this landscape, several actionable principles emerge from the analysis. First, become deeply proficient with AI tools, not just as a passive user but as a skilled practitioner who understands capabilities, limitations, and appropriate use cases. Hassabis's advice to become "unbelievably proficient" applies to all workers, not just young ones; AI literacy is becoming as foundational as computer literacy was in the 1990s. Second, where possible, position yourself within the AI value chain, whether in research, application development, deployment, or governance. Being part of the wave that is reshaping the economy provides both job security and optionality. Third, invest in developing situated expertise and judgment that complements AI rather than competing with it. Domain knowledge—understanding why something works, not just how to execute it—is what allows effective evaluation and direction of AI outputs. Fourth, cultivate skills that involve decision-making under genuine uncertainty, relationship management, and accountability for outcomes; the human-in-the-loop functions that remain essential. Fifth, develop cross-domain knowledge that creates unique combinatorial value: the intersection of AI proficiency with deep domain expertise (medicine, law, engineering, education) is where the highest-value augmentation opportunities lie. Sixth, where possible or relevant, maintain or develop physical-world competencies—whether in healthcare, skilled trades, or other embodied work—which remain protected by AI's current inability to operate in unstructured physical environments.

5. Conclusion

This paper has examined the early evidence on generative AI's impact on U.S. labor markets through the lens of four complementary analytical frameworks, the Anthropic Labor Market Report's novel observed exposure measure, and a range of empirical data sources. The central finding is that while aggregate labor market indicators have remained stable, significant disruption is emerging along specific dimensions: early-career employment in AI-exposed occupations, sectoral contraction in information and professional services, and a nascent shift in the automation-augmentation balance toward greater automation in enterprise applications. These patterns are broadly consistent with the theoretical predictions of the expertise, exposure, and task-specific technical change frameworks, while highlighting the importance of institutional factors emphasized by the occupational power framework and the critical gap between theoretical capability and observed deployment documented by the Labor Market Report.

The evidence supports a cautious but serious assessment of AI's labor market implications. The pace of change is faster than historical precedent for general-purpose technologies, but slower than the most alarming industry forecasts. The distributional effects—concentrated among young workers, in cognitively-intensive sectors, and in geographically concentrated technology hubs. This argues for targeted policy responses rather than either complacency or panic.

Future research should prioritize several directions. First, longitudinal studies tracking individual worker trajectories through the AI transition are needed to distinguish between temporary adjustment costs and permanent displacement. Second, firm-level studies linking AI adoption decisions to employment outcomes would enable sharper causal identification. Third, the interaction between AI and other structural forces including demographic change, climate transition, and geopolitical fragmentation deserves systematic investigation. Fourth, comparative studies across countries with different institutional environments could illuminate how policy and institutional variation mediates AI's labor market effects.

References

Althoff, L., & Reichardt, H. (2026). Task-specific technical change and comparative advantage. Stanford University Working Paper. [Updated March 15, 2026.]

Amodei, D. (2026, January). The adolescence of technology. Retrieved from https://www.darioamodei.com/essay/the-adolescence-of-technology

Appel, R., Massenkoff, M., McCrory, P., McCain, M., Heller, R., Neylon, T., & Tamkin, A. (2026). The Anthropic Economic Index: Economic primitives (v4). Anthropic.

Autor, D., & Thompson, N. (2025). Expertise. NBER Working Paper No. w33941.

Bessen, J. (2016). How computer automation affects occupations: Technology, jobs, and skills. Boston University School of Law, Law and Economics Research Paper No. 15-49.

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.

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.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. OpenAI Working Paper.

Georgieva, K. (2026, January). Remarks at the World Economic Forum Annual Meeting, Davos.

Hassabis, D. (2026, January). Remarks at the World Economic Forum Annual Meeting, Davos.

Huang, J. (2026, January). Remarks at the World Economic Forum Annual Meeting, Davos.

Massenkoff, M., & McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic. Published March 5, 2026. Retrieved from https://www.anthropic.com/research/labor-market-impacts

Microsoft. (2025). New Future of Work Report 2025. Microsoft Research.

Rahman, H. A., Kaynak, E., & Lee, M. T. (2025). An occupational framework for assessing the impact of modern artificial intelligence technologies on work. Northwestern University Working Paper.

Vendraminelli, L., Disorbo, M. D., Hildebrandt, A., McFowland, E., Karunakaran, A., & Bojinov, I. (2025). The GenAI wall effect: Examining the limits to horizontal expertise transfer between occupational insiders and outsiders. Harvard Business School Working Paper 26-011.

World Economic Forum. (2025). Future of Jobs Report 2025. Geneva: World Economic Forum.