

AI & THE FUTURE OF WORK

MARCH 2026

20
26

About Our Company

BCCManagement is Transforming your Business One Digital transformation at a time

BCC Management is a trusted IT business consultancy firm committed to empowering organizations to achieve sustainable growth and operational excellence through innovative, strategic technology solutions. We partner with businesses across industries to align technology initiatives with corporate objectives, ensuring that IT becomes a driver of value rather than just a support function. Our approach is rooted in deep industry knowledge, practical experience, and a clear understanding of the evolving digital landscape.

Our expertise covers a broad spectrum of services, including technology strategy development, IT governance frameworks, IT auditing, IT service management optimization, and comprehensive IT training programs. We also specialize in AI readiness and digital transformation assessments, helping organizations navigate change with confidence. By evaluating existing systems, identifying gaps, and recommending forward thinking solutions, we enable our clients to enhance performance, manage risk, and remain competitive in a rapidly changing environment.

Backed by a proven track record of excellence, BCC Management is dedicated to delivering results that improve efficiency, strengthen compliance, and foster innovation. Our commitment to quality, integrity, and client success defines who we are and drives everything we do.

ASK FOR OUR TRAININGS



AI & THE FUTURE OF WORK

What Data Now Shows About AI's Impact on Employment

Based on Anthropic's Labor Market Impacts Research | March 2026

The gap between what AI can theoretically do and what it actually does in the workplace is vast and that gap is shrinking. This report translates Anthropic's landmark labor market study into strategic implications for business and policy leaders, incorporating new robotics exposure data across 22 occupational categories.

AT A GLANCE

Dimension	Finding	Implication
Displacement risk today	No measurable rise in unemployment for AI-exposed workers	Crisis is not here - yet
Early warning signal	Hiring of workers aged 22–25 in exposed roles fell ~14%	Entry-level pipeline is compressing
Who is exposed	Higher-paid, female, more educated, white-collar workers	AI disrupts knowledge work first
Theoretical vs. actual	AI covers 33% of Computer & Math tasks vs. 94% theoretical max	Enormous runway for further automation
Most at-risk roles	Programmers (75%), Customer Service (72%), Data Entry (67%)	Automation already visible in task data
Robotics complement	Physical sectors (Food, Production, Agriculture) score 50–73% robot potential	AI and robotics target opposite ends of the occupational spectrum

1. CONTEXT: WHY MEASUREMENT MATTERS NOW

Every major technological disruption generates forecasts of mass unemployment and every prior forecast has been at least partially wrong. The automation of routine manufacturing, the rise of offshoring, and the spread of the internet each reshaped labor markets in ways that defied predictions made at the outset.

AI is different enough from prior technologies that extrapolating from historical analogies is hazardous. Yet the research field has struggled to develop frameworks capable of separating early signal from macroeconomic noise. Anthropic's March 2026 study addresses this directly by introducing a new methodology that anchors analysis in real observed usage rather than theoretical capability alone.

The practical stakes are considerable. If displacement arrives gradually, firms and workers have time to adapt. If it is sudden and concentrated, the policy response window is narrow. Understanding which scenario is unfolding and for whom is the central question this analysis addresses.

Why prior measures fell short

Previous frameworks for measuring AI exposure relied on theoretical task assessments: if an LLM could, in principle, perform a task, that task was considered exposed. The Eloundou et al. (2023) measure, for instance, rated 94% of Computer & Math tasks as theoretically feasible for AI.

This approach created a fundamental problem. Theoretical feasibility does not equal economic impact. Legal barriers, workflow integration requirements, human verification steps, and simple inertia mean that large swaths of theoretically exposed work remain untouched in practice. A measure that cannot distinguish between tasks AI could do and tasks AI is actually doing cannot reliably predict employment effects.

Key insight: Theoretical AI capability alone has no statistically significant relationship with projected employment growth. Only when actual usage data is incorporated does the signal emerge.

2. A NEW METHODOLOGY: OBSERVED EXPOSURE

Anthropic's core methodological contribution is the concept of observed exposure a composite measure that combines what AI can theoretically do with what it is actually being asked to do in professional, work-related contexts. This is a meaningful departure from prior research, which treated theoretical capability as a sufficient proxy for labor market risk.

How the measure is constructed

The observed exposure index is built from three data layers, each addressing a distinct limitation of earlier approaches:

Layer	Data Source	What It Captures	Limitation Addressed
1 Task inventory	O*NET database (~800 occupations)	Granular task-level breakdown of each role	Ensures coverage is tied to actual job content, not job titles
2 Theoretical feasibility	Eloundou et al. (2023) β scores	Whether an LLM can, in principle, double task speed	Establishes theoretical ceiling; inputs 0, 0.5, or 1 per task
3 Actual usage	Anthropic Economic Index (Aug + Nov 2025 datasets)	Observed Claude usage in work-related professional contexts	Grounds the measure in behavior, not capability

The three layers are combined through a deliberate weighting scheme. Tasks must clear a minimum usage threshold to be counted as 'covered.' Once covered, they are weighted by the nature of use: fully automated implementations receive full weight, while augmentative use (where a human retains primary authorship) receives half weight. Task-level scores are then averaged to the occupation level, weighted by the time fraction each task represents in the overall role.

The automation vs. augmentation distinction - a critical innovation

The most significant methodological innovation in the Anthropic framework is the explicit differentiation between automated and augmentative AI use. Previous exposure measures treated all AI usage as equivalent, implicitly conflating productivity enhancement with labor displacement.

This conflation matters because the labor market consequences are fundamentally different:

Use Type	Definition	Labor Market Effect	Weight in Index
Automation	AI completes a task with minimal human input; deployed via API pipelines or autonomous workflows	Directly substitutes for labor; reduces headcount need	Full weight (1.0x)
Augmentation	AI assists a human worker; human retains judgment and final output responsibility	May increase productivity per worker; demand effect ambiguous	Half weight (0.5x)

The rationale is straightforward: a customer service platform where AI handles 80% of tickets autonomously creates displacement pressure regardless of whether remaining agents are more productive. An AI writing assistant that helps a lawyer draft briefs faster does not eliminate the lawyer. The index is designed to weight these scenarios appropriately.

Validation and robustness

The methodology includes several validation checks that strengthen confidence in its outputs:

- Alignment with BLS projections: Occupations with higher observed exposure show weaker BLS employment growth projections for 2024–2034, with a statistically significant relationship of -0.6 percentage points per 10 pp of coverage. This correlation does not appear for the theoretical measure alone, validating the incremental power of the usage data.
- Stability across specification choices: The rank-order correlation of occupational exposure is extremely high across many alternative construction choices different β threshold treatments, different usage cutoffs, different automation weighting ratios indicating the measure is not sensitive to arbitrary design decisions.
- Verification against independent sources: 97% of observed Claude usage falls in task categories rated as theoretically feasible by Eloundou et al. ($\beta \geq 0.5$), confirming that real usage is concentrated in theoretically plausible areas rather than being noise.

What the methodology does not capture

Intellectual honesty requires acknowledging the framework's explicit boundaries:

- Single-provider scope: All usage data derives from Claude (Anthropic). Patterns on competing platforms OpenAI GPT-4/o-series, Google Gemini, Microsoft Copilot may differ in task distribution, automation intensity, or occupational focus. The index should be read as a directional indicator, not a census.
- Snapshot in time: Coverage measures are calibrated to Claude capabilities as of late 2025. Rapid capability improvement means the theoretical ceiling itself is moving, potentially compressing the timeline for gap closure.
- US labor market focus: The framework relies on O*NET occupational structures and US employment data. International labor markets with different occupational compositions, collective bargaining arrangements, or AI adoption rates may require adapted measures.
- Lagged wage and hours effects: The primary outcome tracked is unemployment. Early displacement signals that manifest as wage suppression, reduced hours, or shift to gig/contract arrangements are not captured by this version of the framework.

3. A THREE-DIMENSIONAL VIEW OF OCCUPATIONAL EXPOSURE

Extending the Anthropic framework, this analysis incorporates a third dimension: robotics potential. While AI primarily targets cognitive and information-processing tasks, robotics automation targets physical, dexterous, and mobility-dependent work. Together they provide a more complete picture of where automation pressure will fall across the occupational spectrum.

The radar chart below plots all three dimensions simultaneously across 22 occupational categories. The pattern that emerges is striking in its complementarity: AI and robotics target almost opposite ends of the job landscape.

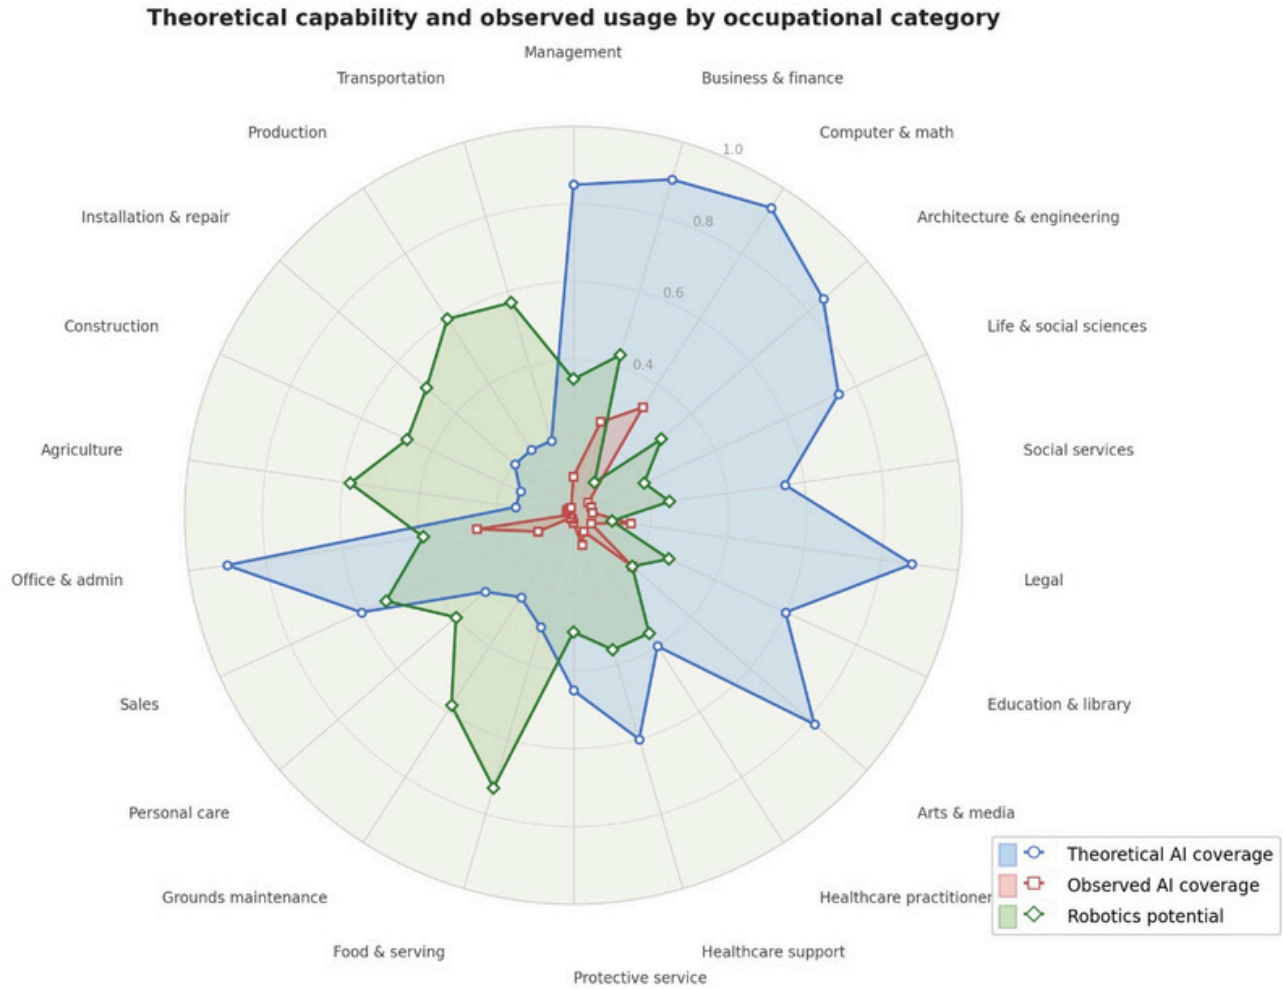


Figure 1: Theoretical AI capability, observed AI coverage, and robotics potential by occupational category. Source: Anthropic Economic Index; Eloundou et al. (2023); MGI robotics analysis.

Reading the chart: three distinct zones

The radar chart reveals three structurally different exposure profiles:

Zone	Occupational Categories	AI Theoretical	AI Observed	Robotics	Strategic Implication
High AI / Low Robot	Management, Business & Finance, Computer & Math, Legal, Office & Admin	85–94%	10–33%	10–43%	Primary AI disruption zone; large gap to close
High Robot / Low AI	Food & Serving, Production, Agriculture, Transportation, Grounds Maint.	15–30%	1–2%	57–73%	Robotics wave; AI has limited near-term role
Zone	Occupational Categories	AI Theoretical	AI Observed	Robotics	Strategic Implication
Moderate Both	Sales, Healthcare, Architecture & Engineering, Installation & Repair	40–85%	5–25%	30–53%	Dual-automation pressure; most complex transition

The complementarity between AI and robotics means total automation pressure is broadly distributed across the occupational spectrum but the timing and form differ sharply. AI disruption is happening now, in knowledge work. Robotics disruption is coming, in physical work.

4. THE EXPOSED WORKFORCE: WHO IS AT RISK

Understanding which workers are most exposed is as important as understanding which tasks are exposed. The demographics of high-exposure occupations defy the popular narrative that AI primarily threatens low-wage, low-skill work.

The most AI-exposed workers are, on average, better paid, more educated, more likely to be female, and more likely to be white or Asian - a demographic profile that distinguishes AI disruption sharply from previous waves of automation.

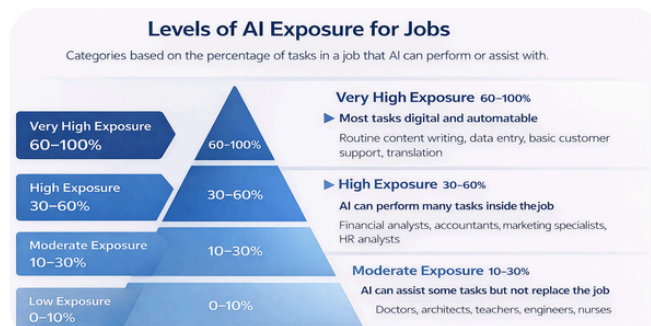
Demographic profile comparison

Characteristic	High Exposure Group (Top Quartile)	Zero Exposure Group	Difference
Female share	Higher	Lower	+16 % points
White share	Higher	Lower	+11 % points
Asian share	~2x more likely	Baseline	~2x
Average wage	Higher	Lower	+47%
Graduate degree holders	17.4%	4.5%	+4x overrepresentation

This shift matters for three reasons. First, high-wage displacement has larger aggregate demand effects than equivalent low-wage displacement. Second, the institutional responses retraining programs, social safety nets, union protections were largely designed for a different kind of disruption. Third, the political economy of AI displacement may differ significantly from historical precedents.

Most exposed occupations

Rank	Occupation	Observed AI Coverage	Primary Exposure Vector
1	Computer Programmers	75%	Code generation, debugging, documentation
2	Customer Service Representatives	72%	Query resolution, API-based automation
3	Data Entry Keyers	67%	Source document processing, data automation
4	Financial Analysts	~55%	Report drafting, data summarization
5	Legal Assistants / Paralegals	~45%	Document review, research, drafting



5. CURRENT EMPLOYMENT EFFECTS: NO CRISIS, BUT EARLY SIGNALS

The most important empirical question is straightforward: is AI causing unemployment? The answer, based on data through early 2026, is: not yet in any measurable aggregate sense but there is a troubling early indicator in hiring data for young workers.

Unemployment: No significant effect

Comparing unemployment trends for workers in the top quartile of AI exposure against those with zero exposure reveals no statistically significant divergence since ChatGPT's release in late 2022. The differential change in unemployment rates is small and indistinguishable from zero. This holds across alternative cutoff definitions, age groups, and when using administrative unemployment insurance claims rather than survey data.

Hiring: A measurable early warning

Workers aged 22–25 entering high-exposure occupations show a statistically significant 14% decline in monthly job-finding rates relative to the pre-ChatGPT baseline. No equivalent decline is observed for workers over 25.

The mechanism appears to be a slowdown in new hiring rather than increased separations firms are not laying off existing staff, but they are reducing intake of new entrants. This distinction has profound long-run implications: workers who cannot enter a field do not accumulate the experience and credentials that enable career advancement.

Outcome	Finding	Confidence	Comparable Study
Overall unemployment	No significant change	High - null result	Consistent with Gimbel et al. (2025)
Young worker hiring (22–25)	~14% decline in job-finding rate	Marginally significant	Consistent with Brynjolfsson et al. (2025)
BLS 10-year projections	–0.6 pp growth per 10 pp exposure	Statistically significant	New finding; no effect for theoretical measure
Unemployment insurance claims	No significant change	High - null result	Consistent with CPS data

Outcome	Finding	Confidence	Comparable Study
Overall unemployment	No significant change	High - null result	Consistent with Gimbel et al. (2025)
Young worker hiring (22–25)	~14% decline in job-finding rate	Marginally significant	Consistent with Brynjolfsson et al. (2025)
BLS 10-year projections	–0.6 pp growth per 10 pp exposure	Statistically significant	New finding; no effect for theoretical measure
Unemployment insurance claims	No significant change	High - null result	Consistent with CPS data

6. STRATEGIC IMPLICATIONS

Translating these findings into organizational and policy action requires confronting a paradox: the absence of current crisis is not evidence that no crisis is coming. The framework is most valuable precisely when effects are ambiguous, which they currently are.

For corporate leaders

01

Audit your exposure portfolio now

Map your workforce against the observed exposure framework. Do not rely on theoretical capability assessments alone, they overstate near-term risk. Focus on roles where automation patterns are already visible in enterprise AI usage data.

02

Protect entry-level pipelines intentionally

The evidence that hiring of young workers is slowing in exposed occupations suggests firms are quietly reducing intake. Firms that continue investing in junior talent will have a structural advantage in three to five years.

03

Reframe AI ROI around automation vs. augmentation

The 47% wage premium among exposed workers means automation delivers outsized cost savings. But augmentation may deliver greater value creation in many contexts. The distinction should drive investment decisions.

04

Monitor the gap, not just current usage

The strategic risk is not today's 33% coverage of Computer & Math tasks, it is the trajectory toward 94%. Build scenario plans around coverage reaching 50%, 70%, and 90% in your highest-exposure functions.

05

Plan for robotics on the physical side

For firms with operations in food service, manufacturing, logistics, or agriculture, robotics exposure is the dominant automation risk. The AI-first planning lens misses this half of the picture entirely.

For policy makers

A

Prioritize labor market entrants in policy design

Existing safety net architecture was designed for displaced workers, not workers who are never hired. The young-worker signal suggests a new challenge: supporting people entering fields that AI is restructuring before they accumulate displacement-triggering work history.

B

Invest in measurement infrastructure

This research exists because Anthropic combined proprietary usage data with public labor statistics. No government agency has equivalent visibility. Public-private data partnerships would dramatically improve the policy response capacity.

C

Revisit education-credential alignment

Graduate degree holders are four times overrepresented in the most-exposed occupational group. If the skills imparted by graduate education are precisely those most replicable by AI, the return on that investment may decline materially.

7. WHAT TO WATCH: LEADING INDICATORS

Indicator	Current Signal	Watch For	Horizon
Young worker hiring rate	~14% decline in exposed occupations	Acceleration or sustained reversal	Quarterly
Coverage gap closure	33% actual vs. 94% theoretical (Comp & Math)	Actual coverage crossing 50%	Annual
Wage premium erosion	+47% premium for exposed workers	Premium compression signaling substitution	Annual
Unemployment divergence	No significant effect to date	Any sustained differential increase	Monthly
API automation share	Growing relative to augmentation	Automation shares exceeding augmentation	Quarterly
Robotics adoption rate	Early physical-sector deployment	Physical-sector job losses in food, production	Annual
Graduate outcomes	Overrepresented in exposed group	ROI decline on AI-adjacent credentials	Annual

APPENDIX: FULL OCCUPATIONAL EXPOSURE REFERENCE DATA

The table below presents the complete dataset underlying the radar chart in Section 3, covering all 22 occupational categories. Data compiled from Anthropic Economic Index usage analysis, Eloundou et al. (2023) theoretical exposure ratings, and MGI robotics potential estimates.

AI Theoretical — Maximum feasible AI task coverage under current LLM architecture	AI Observed — Actual task coverage measured from Claude usage data	Robotics Potential — Estimated physical-task automation potential from robotics	
Occupational Category	AI Theoretical	AI Observed	Robotics Potential
Management	85%	10%	35%
Business & finance	90%	25%	43%
Computer & math	94%	33%	10%
Architecture & engineering	85%	5%	30%
Life & social sciences	75%	5%	20%
Social services	55%	5%	25%
Legal	88%	15%	10%
Education & library	60%	5%	27%
Arts & media	82%	20%	20%
Healthcare practitioners	40%	5%	36%
Healthcare support	60%	8%	36%
Protective service	45%	2%	30%
Food & serving	30%	1%	73%
Grounds maintenance	25%	1%	58%
Personal care	30%	1%	40%
Sales	60%	10%	53%
Office & admin	90%	25%	39%
Agriculture	15%	1%	58%
Construction	15%	2%	47%
Installation & repair	20%	2%	50%
Production	20%	2%	60%
Transportation	20%	2%	57%

Sources: Anthropic Economic Index (Aug–Nov 2025); Eloundou et al. (2023); Massenkoff & McCrory (2026); MGI robotics analysis. AI Theoretical and AI Observed values derived from Anthropic research. Robotics Potential reflects MGI estimates of physical-task automation feasibility.



THANK YOU

Information

This document has been prepared for informational and advisory purposes only. The content reflects the analysis and professional judgment of the authors at the time of preparation and is based on the information available at that time. It should not be considered legal, financial, or regulatory advice. Implementation decisions remain the responsibility of the concerned parties.

While reasonable care has been taken in preparing this document, the authors do not accept liability for decisions made based on its contents or for any direct or indirect consequences arising from its use.

- Anthropic (2026). Labor market impacts of AI: A new measure and early evidence. <https://www.anthropic.com/research/labor-market-impacts>
- Shoushany, R. (2026). AI, robotics, and job displacement: Impacts and projections of the future of work. BCCManagement
- McKinsey Global Institute (2023). Jobs lost, jobs gained
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs

Prepared for:

Rudy Shoushany

BCCManagement Rudy@BccManagement.me