

OpenAI

The AI jobs transition framework

Mapping AI's near-term impact on jobs



Note from the Chief Economist

AI is transforming our lives. It's helping us complete important tasks, reshaping how we work, and expanding what individuals and organizations can do.

Those capabilities also create important questions for the broader economy and job market: what will AI's labor market impact be? Where will the impacts be felt most and when? And how can we ensure that the AI transition works for everyone? While AI capabilities are advancing very quickly, businesses, institutions, and labor markets take time to adjust. That lag means we must avoid two kinds of error: overstating immediate disruption and understating long-run impact.

That is why we see this transition as a force we can shape, not simply observe. Our role is to identify where economic changes are already occurring and disseminate this data widely, so workers, firms, and policymakers can act with the best possible information. We're working to address this gap through sharing what we're seeing and informing a wide set of stakeholders through [releasing data](#), publishing policy ideas like our [Industrial Policy for the Intelligence Age](#) to jump start conversation and debate, and continuing to develop and deploy more advanced models that are easier to use so more people — not just a select few — can participate in the AI economy.

Since I started in this role in 2024, AI capabilities have increased at unprecedented speed, and early adopters are already seeing meaningful productivity gains. Several studies have compared the tasks we do at work to what AI can do today and suggest immediate and widespread economic impact. At the same time, the aggregate labor market impact remains uneven and sometimes contradictory, reflecting the typical frictions that slow technological adoption. But we cannot be complacent. Given how quickly AI capabilities are advancing, there is still potential for broader disruption.

With increasingly capable models on the horizon, we want to help people and organizations understand this transition and plan for it. That includes developing new and more nuanced approaches to estimating the impact of AI on jobs. In this report, we propose a simple framework that moves beyond the traditional AI exposure measures and tests it against what we see in data from the world's largest AI platform, ChatGPT. We combine the following elements: where AI has capability, where humans remain essential, how demand responds to lower costs, and where AI is actually being used today. Together, these factors provide a clearer picture of which jobs are most likely to face pressure first, which are more likely to be redesigned, and which may grow.

We hope this framework contributes to the policy and economic debates about what we need to do to prepare for and navigate the AI transition.

Ronnie Chatterji
OpenAI Chief Economist

Our role is to identify where economic changes are already occurring and disseminate this data widely, so workers, firms, and policymakers can act with the best possible information.

Key takeaways

- 01** AI exposure alone is too blunt a measure to predict which jobs will face near-term disruption.

- 02** Using a new framework across 900+ occupations covering 99.7% of U.S. employment, we find that:
 - 18% of jobs are at a relatively higher short-term automation risk
 - 24% of jobs may see declining employment as task composition shifts while workers remain necessary for key tasks
 - 12% of jobs could grow because of AI as lower effective cost may increase utilization, affordability, access, or quality-adjusted output
 - 46% of jobs are likely to see less change in the short term

- 03** New usage data shows that, among jobs our framework labels as more susceptible to automation, ChatGPT is used about 3x more than in jobs facing less immediate pressure.

- 04** Technical capability does not translate cleanly into job loss risk. Many highly exposed jobs are more likely to be redesigned or expanded than immediately automated.

- 05** These categories are not job-loss forecasts. They are a map for understanding where near-term labor market pressure may emerge first.

Introduction

Most analysis of AI's impact on the labor market begins with the same core question: what jobs are most exposed to AI? That is an important starting point, but it is not enough when it comes to measuring the risk of major labor market disruption. Exposure helps us understand where AI has technical capability. It cannot, on its own, tell us which jobs are most likely to be automated, redesigned, or expanded in the near term.

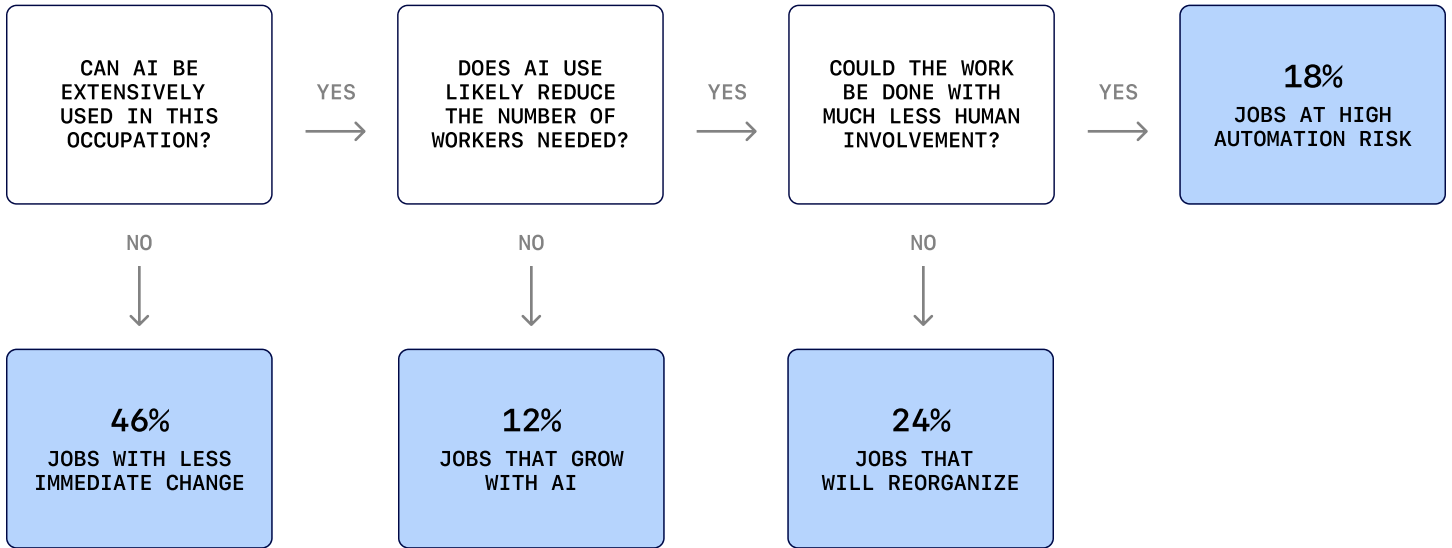
In the three years since a [landmark paper](#) coauthored by OpenAI researchers on AI's potential labor market impacts, AI capabilities have dramatically accelerated, and that trend will only continue as increasingly powerful models are developed. Meanwhile, AI's labor market impacts are uneven and difficult to identify. Even in jobs where AI can do many of the tasks, people are still very much in charge.

This report shows that exposure alone is too blunt to identify where AI may affect jobs first. We introduce the AI Jobs Transition Framework that moves beyond exposure-only measures by combining three dimensions — technical exposure, human necessity, and demand elasticity — and then validating the metric against observed ChatGPT usage data. Across more than 900 occupations covering 153.7 million jobs — 99.7% of U.S. employment — we show that these factors produce a sharper picture of where AI may affect jobs first in the short run. In the long run, many occupation characteristics — and the borders between occupations themselves — are malleable. Additionally, further technological advancements may change the nature of human necessity.

Exposure helps us understand where AI has technical capability. It cannot, on its own, tell us which jobs are most likely to be automated, redesigned, or expanded in the near term.

Using this framework, we sort jobs into four categories and find that:

All 921 occupations (147.9M jobs)



Jobs at high automation risk 18% of jobs are at a higher short-term automation risk

Jobs with less immediate change 46% of jobs are less likely to experience near-term change

Jobs that grow with AI 12% of jobs could grow because of AI

Jobs that will reorganize 24% of jobs may see declining employment as their task composition shifts, but the remaining jobs will still need workers

In other words, many highly exposed jobs are more likely to be reorganized or expanded than immediately automated. Actual AI usage varies widely across these groups. We find that ChatGPT is used about 3x more in the kinds of jobs our framework identifies as most at risk of automation. This is true even when AI can perform similar tasks in less affected jobs. That gap between capabilities and usage makes the limitation of exposure-only measures clear: capability alone doesn't tell us where change in the labor market is actually happening.

Higher exposure doesn't automatically translate into realized usage or short-term job loss risk, nor is more moderate exposure necessarily ironclad protection. An AI-exposed occupation may still remain human-led if regulation, trust, liability, care, or physical-world constraints keep a worker at the center of the service. Occupations like teachers, nurses, and accountants fit this pattern. At the same time, a job with only moderate exposure may still see disruption (1) if the exposed tasks are central, (2) the need for a person to do the job is flexible, and (3) the relatively lower price of output does not increase demand for the goods or services the occupation provides. Some administrative roles fit that pattern better than exposure rankings alone suggest.

We further analyze the relationship between our estimated labor demand elasticity measure and theoretical exposure, showing that there is minimal correlation between the two. This underscores the importance of understanding how the economy might respond to changes in labor productivity and looking beyond occupation characteristics.

These categories are not forecasts of job impacts. Concerns about the potential for broad job displacement remain important. This framework is designed to help identify where those concerns are most likely to emerge in the short term so that policymakers, business leaders, and other institutions can prepare and respond accordingly.

An AI-exposed occupation may still remain human-led if regulation, trust, liability, care, or physical-world constraints keep a worker at the center of the service.

The critical role of people in a world of powerful AI

Many aspects of occupations remain people-led even when AI can do more of the underlying cognitive work.

In many cases, this is a requirement, because people are needed for physical, relational, or regulated tasks. For example, teachers need to be physically in a classroom instructing and caring for students, even while AI can write lesson plans or grade assignments. Court reporters have a very high AI exposure score, ranking in the top 5 in some measures, but in most jurisdictions a human is still legally required to certify each transcript and record. In short, the tasks of court reporters that an AI cannot do, for regulatory reasons, present strong near-term support for continued worker employment in this role. In other occupations, there are less stringent but still important social norms, where people continue to value human interaction even when much of the underlying task could be automated.

Our human necessity measure evaluates the strength of the need for a human worker, asking GPT-5.4 to categorize each occupation. This is the opposite side of the coin of AI exposure, but these two measures capture something slightly different. For software developers, there are two associated O*NET tasks that are theoretically less exposed to AI: training users and meeting with management. These are now things AI can do in some formats. On the other hand, while occupations in the legal industry have relatively high exposure scores, there is a regulatory bar prohibiting AI from doing many parts of the job, ensuring the need for workers to perform many functions.

We match occupations with human necessity into one of three categories:

Regulatory and accountability necessity applies when the job requires a licensed person to approve or make decisions, or when a person still has to take responsibility for the result.

Relational necessity applies when the value of the service depends on trust, care, persuasion, teaching, or human connection. This is largely based on what people prefer today and could change over time.

Physical necessity applies when the occupation requires physical action in the real world, including in-person care, field inspection, or physical execution.

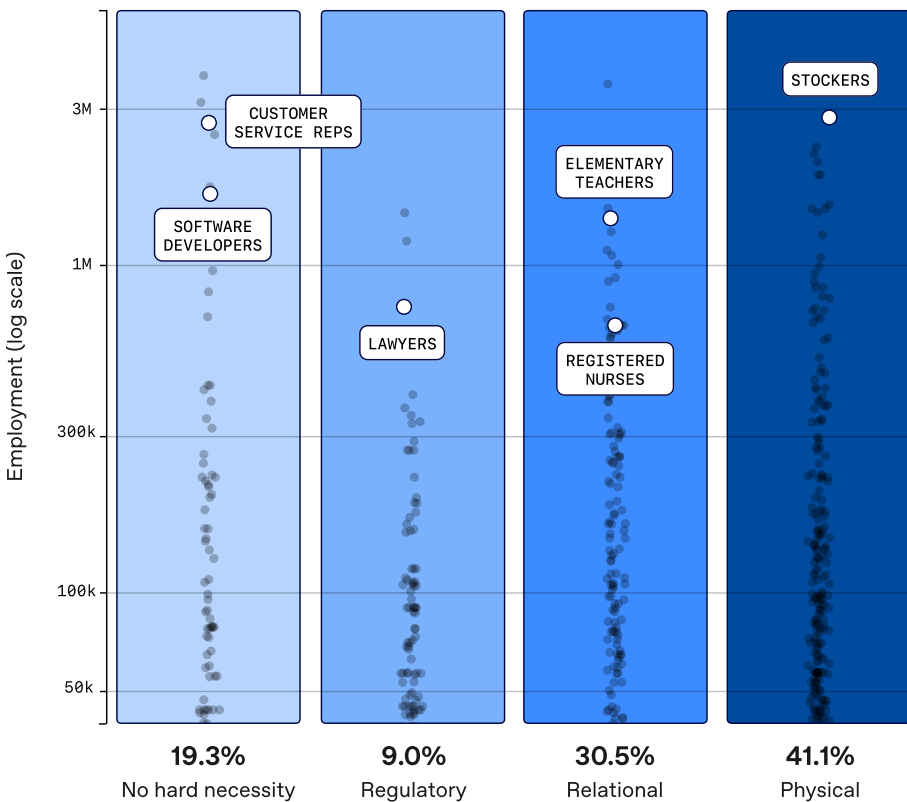
Measuring the importance of workers helps explain why exposure does not map one-to-one onto automation risk. A teacher may face high exposure in lesson planning or content generation, but relational and institutional necessity remain strong: teachers must be physically present to instruct students. A nurse may face high exposure in documentation and information synthesis, but physical and accountability necessity remain central: nurses must care for patients directly in the clinic and at the bedside.

However, people remaining in the loop is not the same thing as an unchanged labor market. A lawyer may remain necessary because of accountability and representation, but if each lawyer can handle much more work and lower prices don't increase demand very much, employment can still decrease. Even in jobs with fewer tasks that AI can do directly, work may transform as AI expertise allows less experienced workers to do more complex problem solving in the field. A high-human necessity occupation can still face labor-reducing pressure if AI raises worker productivity by changing the task composition of the occupation and demand for the good or service does not expand enough to absorb it. These necessities may also change over time either as technology advances (e.g., progress on robotics could change physical necessity) or as societal preferences evolve.

Figure 1 classifies occupations by "Human Necessity," which indicates mandatory human involvement due to physical, relational, or accountability constraints, moving beyond AI's cognitive capabilities. These categories — Regulatory/Accountability (e.g., judges, lawyers), Relational (e.g., teachers, nurses), and Physical (e.g., plumbers, physical therapists) — illustrate why many highly-exposed jobs will be augmented, not automated.¹ High AI capability is often gated by the need for a licensed signature, fiduciary duty, or in-person emotional interaction, preventing full substitution, at least in the short run. This does not include broader preferences for human interaction, which would include occupations like customer service representatives. While AI may already be able to do almost all, if not all, of the tasks of a virtual customer service representative, many humans continue to prefer to speak to other people, which has important implications for the evolution of the labor market. We use the stricter version in what follows, but want to be clear that it doesn't fully capture this human preference component.

Fig. 1

Human necessity helps us understand where workers remain vital



¹ In the case that more than one classification is valid, we prioritize physical, regulatory, and then relational necessity.

Why higher worker productivity doesn't always mean fewer jobs

Demand elasticity — how much demand changes when price changes — is what connects productivity to employment.

If AI makes the cost of providing a good or service cheaper, the effect on employment in related occupations is ambiguous. Simply, when goods become cheaper, people often buy more of them, sometimes leading to an increase in employment in affected sectors, often referred to as Jevon's Paradox.

In the spirit of Asirvatham et al. (2026), we use GPT-5.4 to estimate the potential response to a productivity shock of demand for workers in a given occupation, based on GPT's deduced understanding of latent labor market structure. We depict the distribution of occupations by estimated demand response in Figure 2. We detail the estimation methodology in the Appendix.

For each occupation, we construct a profile using the O*NET occupation title, summary description, task statements, detailed work activities, intermediate work activities, and work activity elements. We then ask GPT-5.4 to identify the relevant output and estimate the quantity response to a standardized 10 percent decline in the relevant price of that output over a two-to-three-year horizon. The prompt directs the model to distinguish occupations where lower prices could unlock latent demand, greater frequency of use, more customization, or new customers. These estimates are rough indicators of whether lower lower costs will lead to more demand for a job's output, not precise causal measurements/predictors.

Why higher worker productivity doesn't always mean fewer jobs

Breaking down occupations by estimated demand response:

The least elastic occupations include firefighters and home health aides.

Somewhat elastic occupations include physical therapists, editors, and dental hygienists.

Some of the **most elastic** occupations include graphic designers and software developers.

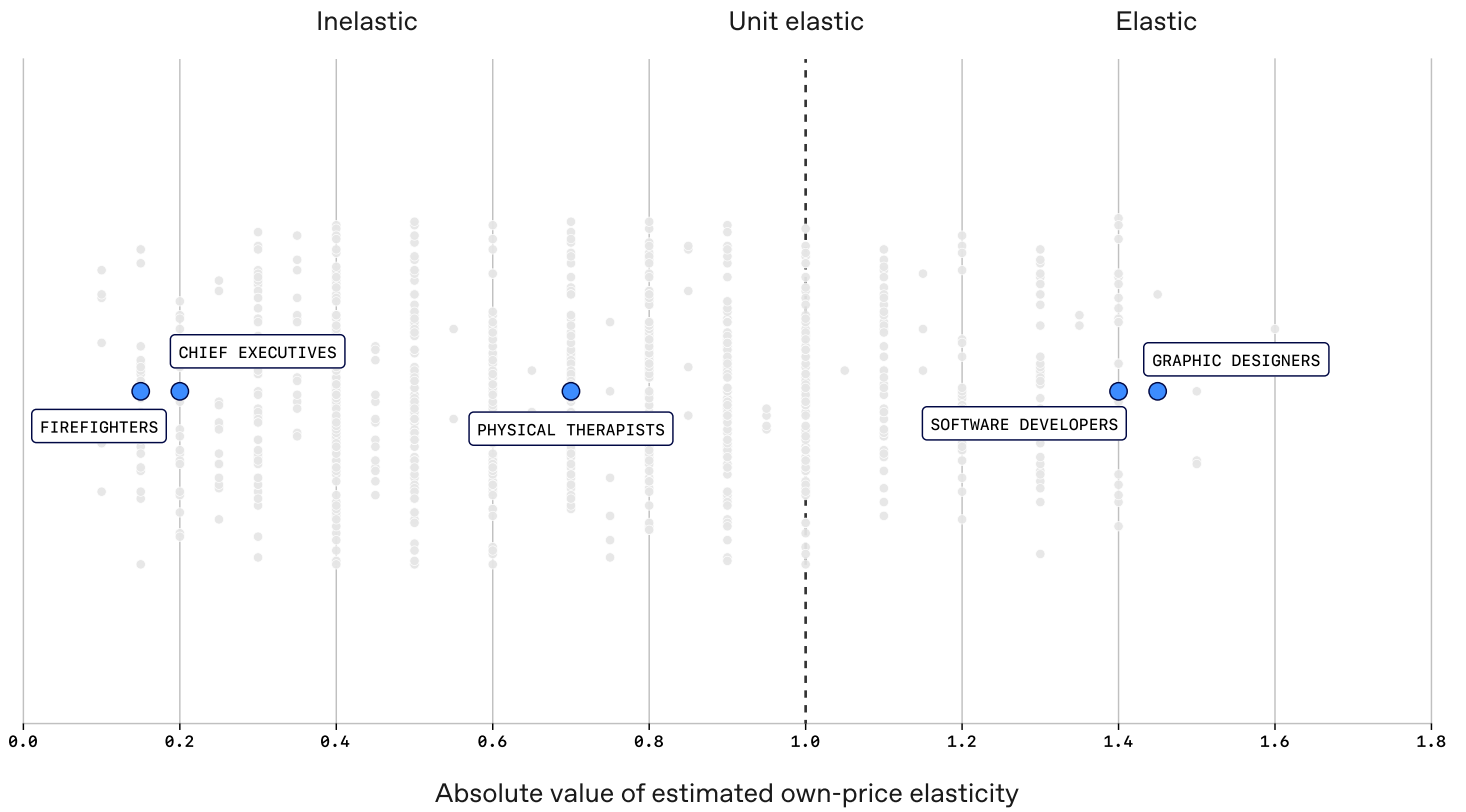
These estimates are largely consistent with those available for comparison in the economics literature. For example, Autor and Dorn (2013) note that demand for manual, task-intensive work is often relatively price inelastic; we see this in our estimates as well, with Home Health Aides at -0.8, Childcare Workers at -0.7, and Maids and Housekeeping Cleaners at -0.8. Angerhofer et al. (2025) find a labor supply demand elasticity for teachers in Wisconsin of -0.72; our estimates for teachers range from -0.3 to -0.45.

We also note that there could be a point where AI automation reduces the costs of tasks so dramatically and makes their output so abundant, that it overwhelms demand elasticity and these occupations see more overall automation and decreasing employment. We see this as being a potentially critical part of the long-run effects of AI on the labor market, but it is unclear how quickly that will unfold.

Why higher worker productivity doesn't always mean fewer jobs

Fig. 2 Occupation by GPT-estimated demand elasticity

Many occupations are estimated to have modest demand expansion when output gets cheaper



AI jobs transition framework

When combined, these measures help progress the conversation from technical possibility to labor-market relevance.

The framework asks four questions:

01 Can AI do a meaningful share of the work?

02 If AI lowers the effective cost of providing the service, is demand likely to expand enough to absorb the productivity gain?

02 If it can, for remaining tasks, is a person still central to the work's delivery, judgment, accountability, or physical execution?

04 Is AI already being used meaningfully for these tasks?

The examples below illustrate how these dynamics play out in different cases:

Weak human necessity and low demand elasticity may create direct automation pressure.

Strong human necessity may preserve the need for a worker but reduce staffing needs by increasing worker productivity.

Lower cost may expand access, utilization, or service quality enough to support continued or even higher employment.

Using this approach, we group occupations into four different possible outcomes in the AI transition:

Jobs at higher automation risk have high exposure, weak human necessity, and insufficient or ambiguous demand offset.

Jobs that will reorganize have high exposure and strong human necessity, but demand is not elastic enough to absorb the productivity gains.


Jobs that grow with AI have high exposure and enough demand response that lower effective cost may increase utilization, affordability, access, or quality-adjusted output.

Jobs with less immediate change are those where the current combination of exposure, necessity, and elasticity does not yet point clearly to one dominant near-term outcome.

Overall, we find that about 46% see less immediate change, 25% will reorganize, 18% are at higher automation risk, and 12% grow with AI. This is mapped in Figure 3.

Fig. 3

Employment share by archetype



These shares should be read as a map of labor-market transition risk, not precise forecasts of displacement. They indicate where the current combination of capability, adoption, human necessity, and demand offset points most clearly toward a particular type of adjustment pressure.

The value of this framework is that it recovers differences that are invisible in exposure-only measurements, in which the nuance between jobs that are completely automatable, jobs which may very well grow, and jobs that still require workers is lost. Even within the most exposed occupations, jobs vary in both human necessity and demand elasticity. As a result, many highly exposed jobs are better understood to be opportunities for redesign and scale-expansion rather than immediate automation, while others are better understood as future-risk where adoption has not yet caught up with technical possibility.


Capability overhang by job

This section builds on the theoretical framework by validating it against real-world AI usage data, confirming that current and theoretical AI usage is consistent with these broad categorizations. Recent work by our team on the capability overhang, as well as work by Massenkoff and McCrory (2026), highlights how actual AI usage lags AI capabilities. We confirm this finding at the occupation level by relying on proprietary data on aggregated and anonymized work-related ChatGPT usage for consumers by task.

We map ChatGPT usage into O*NET tasks, and then aggregate them into occupations based on an estimate of how much time an occupation spends doing that task, comparing the observed share to the theoretically expected share if ChatGPT was used for the full range of the task in this role. We do not observe or infer the occupations of any user, relying on the assignment of O*NET tasks to occupations. We compare this to the Eloundou et al (2023) measure of what AI can do. This lets us estimate observed exposure: the share of a job's time spent on tasks where AI is actually being used. We estimate this by comparing the usage share to the share we would expect for a given task, given the estimated man hours performed on this task in the broader U.S. economy.

Fig. 4

The capability overhang by occupation



In Figure 4, across every job category, current usage lags behind the possible. Some fields are seeing more and more AI use — particularly legal, education, and office/administrative jobs. Some fields — like food prep and serving, those where AI’s current capabilities are less relevant to the job — see less usage, even for tasks that AI can do. For theoretical exposure, the mean score is derived from Eloundou et al. (2023). Further details can be found in Appendix 2.

Translating the framework into the real world

We then use our estimates of the theoretical and realized exposure to understand how consumer ChatGPT usage distinguishes each archetype. Figure 5 shows the capability overhang — the gap between AI's theoretical potential and its observed use — across job archetypes, confirming AI's capabilities far exceed current realization. Both the realized usage and the capability overhang are highest for jobs that our framework identifies as being at the highest automation risk. This overhang is uneven: jobs at higher automation risk show the highest adoption but also the largest gap between realized and theoretical exposure. This substantial difference for jobs that will reorganize, grow with AI, or face less change indicates AI adoption is limited not just by technical feasibility but by human necessity, demand, and institutional friction. Thus, exposure alone is a weak predictor of immediate labor market pressure, validating the transition framework's core thesis.

Fig. 5

Capability overhang by archetype

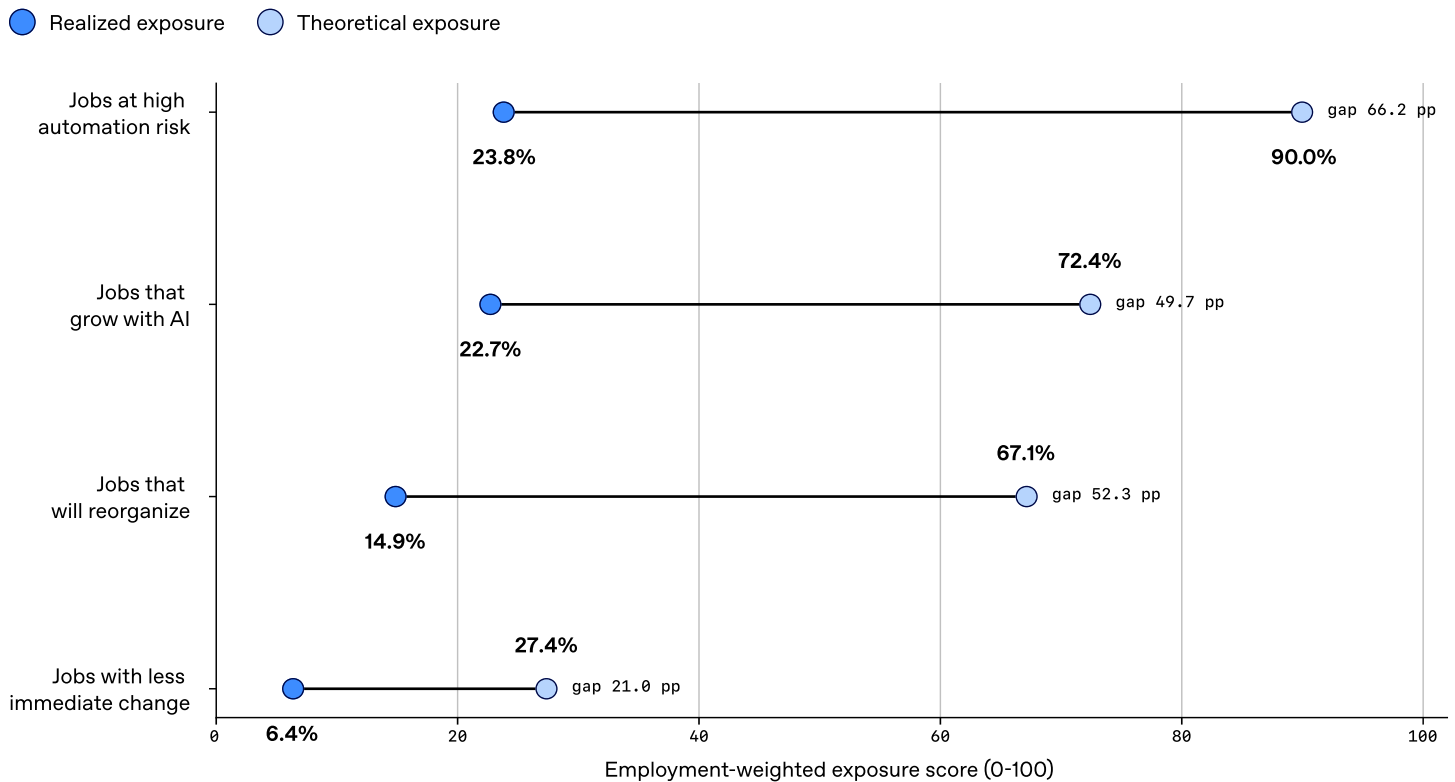
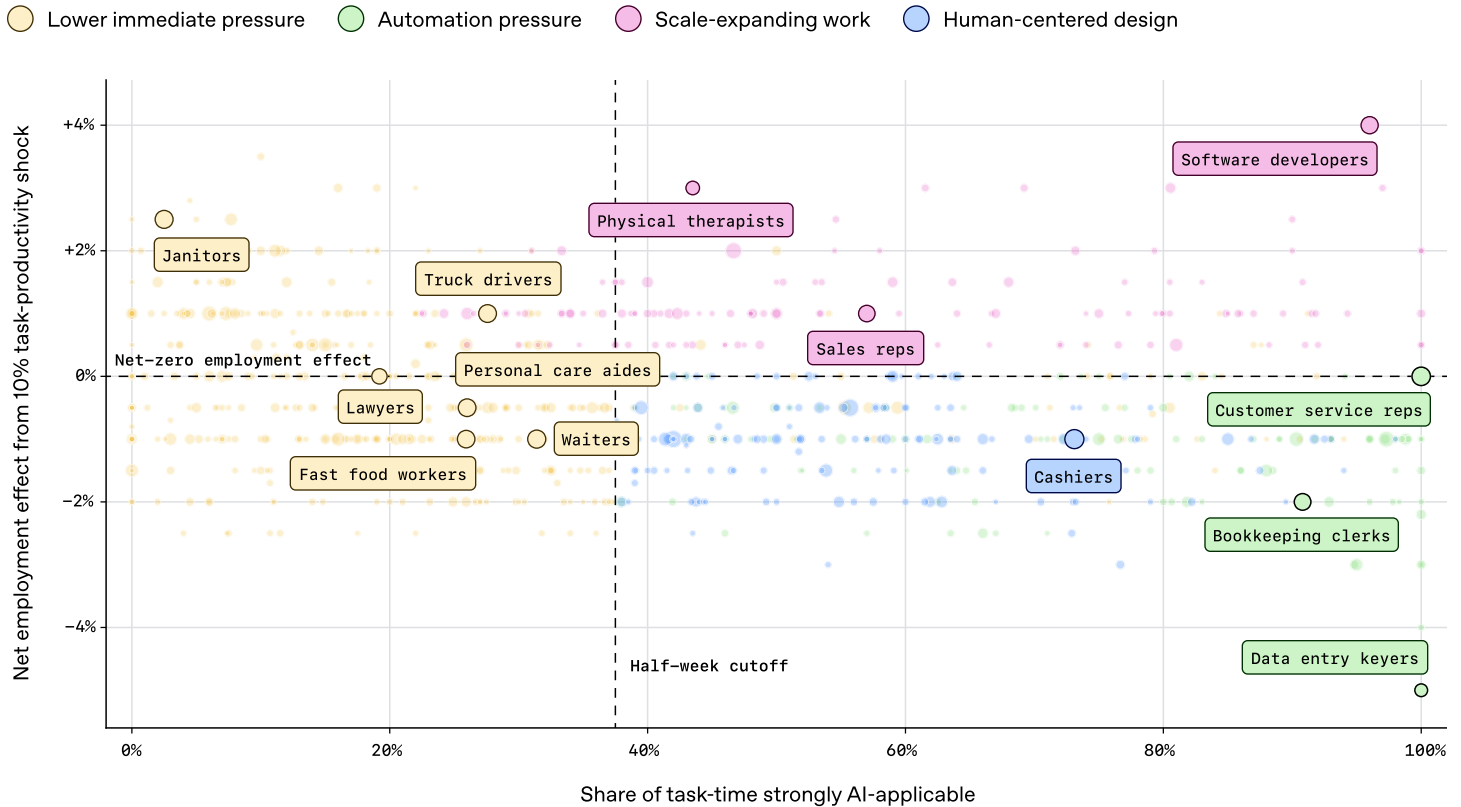


Fig. 6

High exposure occupations split by whether scale effects offset direct labor savings

Each dot is an O*NET occupation; the vertical line marks where strongly AI-applicable tasks make up roughly 15 hours of a 40-hour week



Bubble size reflects 2024 employment. Net employment effect = direct labor-saving effect + scale effect from expanded output for a 10% effective task-productivity shock.

Figure 6 further illustrates why broad exposure rankings are too blunt for policy: AI exposure is not tightly correlated to how labor demand may respond to changes in labor productivity. What matters is not simply whether an occupation is exposed, but where it falls in the transition framework. There are many occupations in the top right quadrant of the chart: these are jobs which an exposure measure would suggest are highly exposed to AI but who could see a demand increase in response to AI-related productivity enhancements. Occupations with higher demand elasticity are the clearest candidates for scale effects, where AI may expand access or output rather than primarily reduce labor demand.

Demand elasticity provides the demand-side input behind the net employment response shown in Figure 6. A higher absolute own-price elasticity means that an AI-driven reduction in effective cost is more likely to expand output enough to offset direct labor savings, moving occupations toward the positive side of the net employment axis. Low-elasticity occupations can still be highly exposed, but in the scatter they remain closer to the labor-saving side because demand is less likely to grow enough to absorb the productivity gain.

The jobs at higher automation risk in the lower-right portion of the figure are the clearest candidates for early-warning monitoring, adjustment support, and localized transition planning; these are the occupations where observed AI exposure should rightfully align with intense focus from policymakers. Many of the occupations in this quadrant are those that could reorganize: they have a strong requirement for workers, but they could see employment declines as the task content of the occupation shifts to focus on the human-centered tasks. Policymakers should also carefully consider these occupations in crafting potential policy responses.

Next, we aim to understand whether or not these occupation differences are already showing up in the labor market. By occupation group, since the first quarter of 2024, people that are or were employed in jobs we predict to have less immediate change have actually seen the largest increase in unemployment at +0.6 p.p.² For comparison, people employed in jobs at higher automation risk or jobs that will reorganize have seen +0.3 p.p. and +0.3 p.p. increases in the unemployment rate over the same period, respectively.

This is why, while AI may be linked to some changes in employment, it remains difficult to clearly link AI and the aggregate labor market. This could change rapidly, which is why these occupations warrant continued close monitoring. There is also a broader implication: governments should invest in better occupational measurement, so that the public can better understand, evaluate, and respond to the effects of AI in the labor market.

² We compute the unemployment rate using CPS Basic Monthly labor-force records. Workers are assigned to one of the report's occupation archetypes using their current occupation if employed, and We compute the unemployment rate using CPS Basic Monthly labor-force records. Workers are assigned to one of the report's occupation archetypes using their current occupation if employed, and their most recent occupation if unemployed. CPS occupation codes are bridged to SO%*NET occupations, then allocated to the report's archetypes using the same occupation spine as the main figures. For each archetype and month, we sum CPS-weighted employed workers and CPS-weighted unemployed workers to compute the unemployment rate.

Table 1

Recent unemployment rates by archetype

	2024 Q1 Unemployment rate	2026 Q1 Unemployment rate	Change
Jobs at high automation risk	3.3%	3.6%	+0.3 pp
Jobs that will reorganize	2.9%	3.2%	+0.3 pp
Jobs that grow with AI	3.5%	3.8%	+0.3 pp
Jobs with less immediate change	4.5%	5.1%	+0.6 pp

Policy implications by job archetype

These archetypes point to different policy responses based on the kinds of challenges and opportunities workers in these roles may face in the near term. These ideas should be considered alongside those contained in [OpenAI's Industrial Policy for the Intelligence Age](#).

For jobs with higher automation risk, the priority is helping workers navigate change. These jobs are the strongest candidates for early warning systems, transition assistance, reskilling support, and local monitoring of labor-market disruption, especially when that support can help workers move into roles where people remain essential.

For jobs that will reorganize, the priority is shaping how AI changes work, not just how much work gets done. People will still be needed for important parts of these jobs, but changing task mix and higher productivity could still reduce headcount in some roles. Policymakers should focus on staffing standards, professional guidance, and other guardrails that shape workload, service quality, and worker autonomy.

For jobs that grow with AI, the priority is helping more workers and communities benefit from that growth. These are the clearest cases for capacity-building, procurement reform, reimbursement redesign, and other policies that expand access to these jobs and the services they support.

For jobs with less immediate change, the priority is better measurement and continued monitoring. These occupations should not be treated as permanently insulated, but they are less likely to be the first places where AI-driven labor market pressure shows up.

Caveats and interpretation

A critical limitation of this framework is that it simplifies how tasks are actually organized within occupations.

Recent work underscores that the location of bottlenecks and complementarities within the workflow is critical for understanding automation. In practice, automating some tasks may increase the value of the remaining human tasks by freeing time and concentrating effort on the parts of the job that are hardest to remove. As a result, this framework may still overstate substitution risk, understate augmentation and redesign, and miss the possibility that adoption unfolds in threshold-like or bundled ways rather than smoothly. For that reason, these results should be read as a policy-oriented map of where transition pressure may emerge and what form it may take, not as a precise forecast of near-term job loss.

Elasticity is the least directly observed object in the current framework. The values are GPT-generated proxies relying on the latent structure of the labor market rather than estimated from clean occupation-level demand systems, and they should be interpreted as structured approximations of how output demand is likely to respond to lower effective cost, not as a substitute for high-quality causal evidence.

Many of these occupation characteristics may also change over time. AI may impact not only wages and the price of work, but also quality, speed, convenience, customization, and access, all of which can shift demand in ways this framework only partly captures. The framework also does not fully capture general-equilibrium wage effects, cross-occupation spillovers, induced demand in complementary sectors, or macroeconomic feedback through income and prices.

The big picture

Exposure is only the starting point. A more actionable signal comes from how it combines with human necessity and demand response. Those interactions shape whether AI reduces labor demand, changes how work is done, or expands the scale of economic activity.

In this framework, occupation exposure helps us identify which occupations overlap with AI capabilities, but it is not presently strongly related to employment changes. The gap between actual and theoretical exposure distinguishes genuine insulation from delayed diffusion. Human necessity and demand elasticity help explain why some occupations may remain human-led, why others may see staffing compress, and why others may expand as lower costs increase output.

Our ability to forecast far into the future is limited, and it is very difficult to project how much the labor market might evolve over the long term. On a shorter horizon, this framework should help us envision how the labor market may evolve and what policy responses we can consider and implement to facilitate a smoother and more people-centered AI transition.

Appendix 1: Methodology

We build four occupation-level measures, each meant to capture a different mechanism through which AI may affect labor demand. The central idea is that “exposure” alone is too narrow. We therefore measure not only where AI can affect the work, but also where it is already being used, where a human remains hard to remove, and where higher productivity is likely to reduce employment versus expand output.

Theoretical exposure starts from a task-level assessment of which work activities are technically exposed to large language models from Eloundou et al. (2023). A task is counted as theoretically exposed if prior research judged that the task could be affected by LLMs, either with additional tools or directly through the model itself.

Our observed exposure measure is constructed from aggregated and anonymized work-related consumer ChatGPT usage data and linked back to occupations through ONET task structure. This approach is further detailed in Appendix 2.

Our human necessity measure is built from a strict occupation-level taxonomy of hard human bottlenecks. Occupations are classified by applying structured rules to occupational titles and descriptions, with the goal of identifying cases where the worker remains difficult to remove even if AI can do more of the underlying cognitive work. The strict version focuses on live care and teaching, licensed or accountable judgment, and physical execution in the real world.

Our demand elasticity measure is also estimated at the occupation level from structured rules applied to occupation titles and descriptions. We estimate elasticity depending on whether the occupation’s output appears to be essential and price-insensitive, derived from fixed budgets, or more discretionary and responsive to lower cost. Occupations in health, care, regulation, and core public services tend to receive lower elasticity estimates; occupations in sales, marketing, travel, hospitality, entertainment, and other discretionary services tend to receive higher ones.

Finally, we combine these raw measures into derived labor-market transition measures. Observed exposure defines the scale of possible change. Human necessity helps separate cases where AI can substitute directly from cases where a person remains central to delivery. Demand elasticity helps separate cases where productivity gains mainly reduce labor needs from cases where lower cost may expand output. Those combined measures are then used to sort occupations into our labor-market archetypes. Further methodological details and robustness exercises will follow in a technical appendix.

Appendix 2: Mapping AI usage to occupations

This analysis compares two concepts: theoretical AI exposure and realized AI exposure. Theoretical exposure measures how much of an occupation's work could potentially be affected by large language models, based on task-level assessments from prior research. Realized exposure measures how much of that potentially exposed work appears in observed work-related ChatGPT usage. The analysis covers work-related ChatGPT activity from the second half of 2025.

Theoretical exposure

Theoretical exposure starts from a task-level assessment of which work activities are technically exposed to large language models from Eloundou et al. (2023). A task is counted as theoretically exposed if prior research judged that the task could be affected by LLMs, either with additional tools or directly through the model itself.

Each occupation is represented as a bundle of tasks. Because not all tasks take the same amount of time, tasks are weighted by their estimated importance in the occupation's overall workload. The theoretical exposure score for an occupation is the share of its task time spent on tasks that are theoretically exposed to AI.

For example, if 70 percent of an occupation's task time is spent on tasks judged to be AI-exposed, that occupation's theoretical exposure is 70 percent.

Realized exposure

Realized exposure asks a different question: of the task time that is theoretically exposed, how much is already reflected in observed work-related ChatGPT use?

To answer this, aggregated and anonymized work-related consumer ChatGPT conversations are linked to occupational task categories. The analysis focuses only on conversations identified as work-related. Each linked task category is then compared with how common that task would be expected to be in the labor market, based on occupational task weights and employment.

This produces a measure of usage intensity: whether a task appears in ChatGPT use more or less often than expected given how much labor time that task represents. A task that appears at its expected rate is treated as fully realized. A task that appears at half its expected rate is treated as half realized. A task that appears more often than expected is capped at fully realized, so that no task can count for more than its share of occupational task time. The realized exposure score for an occupation is then calculated as the share of total occupational task time that is both theoretically exposed and observed in work-related ChatGPT usage, after applying this cap. We will index to this level moving forward so updates to this measure will be able to identify relative growth in realized exposure.

Aggregation across occupations

The occupation-level scores are also aggregated into broad occupational groups. These group-level results are weighted by employment, so larger occupations contribute more to the group average than smaller occupations.

Appendix 3: Elasticity estimation

We estimate own-price elasticity at the occupation level by first defining what each occupation actually produces. For every ONET-SOC occupation, we build a structured profile from ONET data: the occupation title and description, core and supplemental tasks, detailed work activities, intermediate work activities, and broader work-activity elements. This gives the model a consistent description of the “output” of the occupation, whether that output is legal advice, software features, restaurant service, medical care, management oversight, public safety, or administrative support.

We then ask GPT-5.4-mini to evaluate a standardized price-change scenario for each occupation: suppose the customer price of the occupation’s output falls by 10% because production costs fall, while quality and product scope stay the same. The model is instructed to estimate only the demand response for the occupation’s goods or services, not AI exposure and not the demand for workers themselves. In plain terms, the question is: if this occupation’s output became cheaper to buy, how much more of it would people, firms, governments, or institutions want over the next two to three years?

For each occupation, the model returns a central own-price elasticity estimate using the standard economics convention: elasticity equals the percent change in quantity demanded divided by the percent change in price, so the number is normally negative. For example, if a 10% price decrease would increase quantity demanded by 15%, the elasticity is -1.5.

The prompt explicitly pushes the model to distinguish occupations where demand is genuinely price-rationed from occupations where demand is mostly fixed by other forces. Elasticity is made more negative when lower prices would unlock latent demand, new buyers, more frequent use, or more customization. It is kept closer to zero when the quantity demanded is constrained by emergencies, legal mandates, fixed staffing structures, public budgets, insurance reimbursement, licensing, physical incidence, or narrow unavoidable caseloads. We treat the resulting estimate as a “price-response prior”: a disciplined, occupation-specific estimate of how demand for the occupation’s output responds to price, not a full forecast of employment effects from AI.

Demand elasticity is the demand-side bridge between the occupation-level productivity shock and the net employment response shown in Figure 6. The own-price elasticity estimate should be read as a structured prior on how much additional output demand might be induced when AI lowers the effective cost of producing an occupation’s output. The net employment response is not the elasticity estimate itself; it combines the direct labor-saving effect of higher task productivity with the scale effect from expanded output. Occupations with higher absolute own-price elasticity are therefore more likely to have positive or less negative net employment responses, because lower effective prices, shorter wait times, improved quality, or expanded access can increase output volume enough to offset some of the reduction in labor required per unit.

Appendix 3: Elasticity estimation

By contrast, low-elasticity occupations may still have high AI exposure, but the scatter places them closer to the negative side of the net employment axis because demand is less likely to expand enough to absorb the productivity gain. Thus, Figure 6 should be interpreted as a transition-space map: exposure identifies where AI can affect tasks, while elasticity helps explain whether those productivity gains translate mainly into employment compression or into scale expansion.

Suggested citation

Richmond, Alex Martin. 2026. *The AI Jobs Transition Framework: Mapping AI's Near-Term Impact on Jobs*. OpenAI Economic Research.