OpenAI

Work at the Frontier:

How workers are unlocking new ways of working

September 2026

Alex Martin Richmond and Caroline Chin

This is the second Work at the Frontier report from the OpenAI Economic Research Team. This new paper, building off the first Work at the Frontier report published in July 2026, continues our series exploring how work is changing across the economy. The first report showed that workers commonly use AI to explore or take on tasks outside their traditional job description. This report asks what happens next: how workers actually do the new work beyond their normal scope, and whether these efforts become part of their ongoing regular work activity.

Key takeaways

- **Workers use AI differently for tasks outside their own jobs.** Compared to tasks associated with their current jobs, when workers take on new tasks they typically use shorter prompts and are less likely to ask for explanations, a particular format, or advice, indicating that workers could be relying more on AI expertise to carry out these types of tasks.

- **These new tasks become a regular part of how workers use AI.** Our research shows that workers return to tasks outside of their own jobs more frequently, and these tasks take up a larger share of their overall AI work over time, suggesting that this usage may represent a more fundamental shift in how people work than brief periods of experimentation.

- **Some of these new tasks are stickier than others.** New tasks like customer communication and advertising or promotional writing are more likely to get adopted into workflows than legal research and explaining financial information to customers, showing that new tasks are not incorporated uniformly into work routines.

Introduction

The rise of the modern corporation brought the creation of specialized departments such as finance, legal, marketing, and human resources. This form of specialization often means that completing a task requires several people working together and the tightly defined functional roles. The emergence of generative AI represents a new opportunity. By helping workers handle more work themselves, AI could make it easier to complete complex projects.

As workers increasingly use AI tools, they seek help with tasks both within and beyond the boundaries of their occupations. Much of their work involves what we call **“within-occupation” work**. In such situations, workers turn to AI for help with tasks directly linked to, or closely related to, their identified roles. This could include marketing specialists seeking input on social media initiatives or software engineers seeking help writing code. At the same time, today’s workers often embark on what we call **“cross-occupation” work**, too. In such cases, AI helps them pursue tasks outside the work within their traditional occupation or occupational boundary—a pattern we call **“task crossover”**.

Our first “Work at the Frontier” report details this pattern, highlighting workers’ growing use of AI for tasks outside their traditional roles. This report, the second in the series, goes further to understand how workers use AI for tasks outside their jobs and how this pattern changes over time. To do this, we study an expanded sample of more than 1.5 million work-related ChatGPT messages from individual accounts sent by U.S.-registered users during a four-month period from April 2026 through July 2026, identifying users’ occupations from role or department information provided during their ChatGPT Business onboarding.[1]

We find that workers use AI differently for tasks outside their occupations, and that some of these activities become recurring parts of their AI use. Together, these patterns suggest one way jobs could broaden: through changes in the mix of activities workers undertake.

Task crossover: implications for AI usage

Imagine how a worker might write a prompt for a task with which they are highly familiar. They are experts in their field and know with a degree of certainty what they are asking AI to provide. In contrast, imagine what they might do differently if they were approaching a wholly new area of engagement beyond their formal, professional expertise. In our research, we explore this level of engagement in a large-scale analysis of more than 1.5 million classified messages sent by OpenAI users in our four-month study of their activities from April 2026 through July 2026.

In general, cross-occupation requests are noticeably less likely to ask for explanations or “how-to” help. The cross-occupation requests also are less likely to specify a format or ask for advice. At the same time, cross-occupation requests are more likely to provide examples or background; they also are more likely to involve checking or verification. In addition, cross-occupation prompts are slightly shorter than within-occupation prompts, by about 20 characters on average.

[1] We exclude training-disabled data and messages without usable classifications. We carefully preserve user privacy, aggregating and anonymizing messages for analysis. In doing this work, researchers on our team never read individual user messages.

Figure 1

Cross-occupation prompts are less likely to request explanations or how-to help

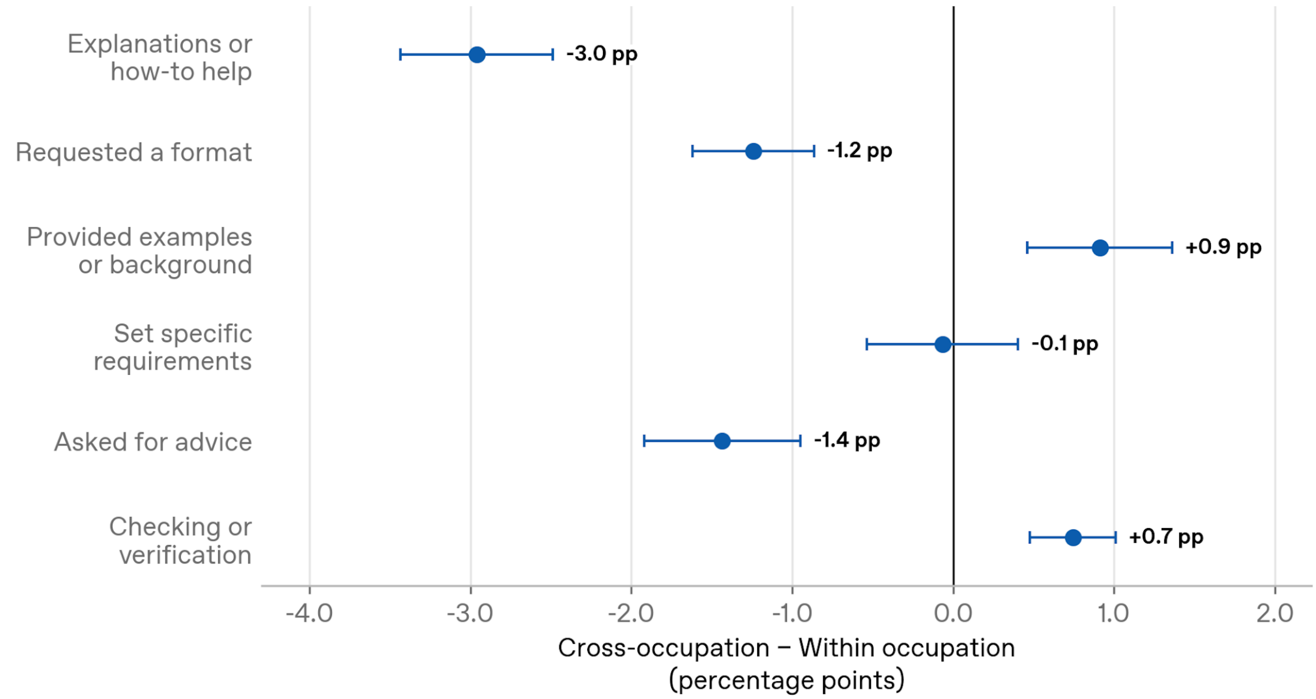

Note: In this chart, negative values indicate lower values for cross-occupation work relative to within-occupation work. We first summarize each worker's messages in the two categories separately and calculate the difference between those summaries. We then average these differences across workers, giving each worker equal weight. Bars show 95% confidence intervals.

We classify prompt content into the following categories, which can overlap:

- **Explanation or how-to requests** include walkthroughs, step-by-step help, templates, and checklists.
- **Requesting a format** means asking for a specific structure or output type.
- **Providing examples or background** means a user supplying examples, attachments, excerpts, or other context.
- **Setting specific requirements** includes audience, tone, length, and what to include or avoid.
- **Asking for advice** combines broad suggestions and specific recommendations.
- **Checking or verification** includes checking accuracy, asking for sources, and assessing assumptions or risks.

Figure 2

Cross-occupation prompts are shorter

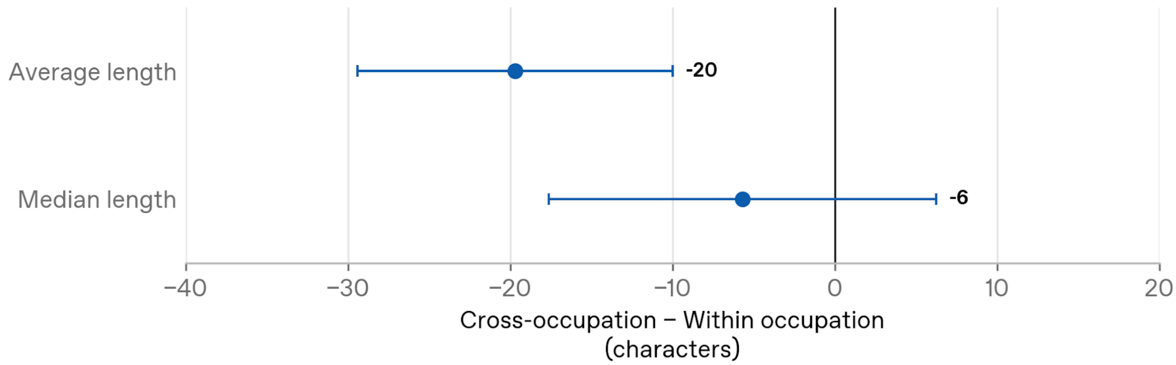

Note: In this chart, negative values indicate lower values for cross-occupation work relative to within-occupation work. We first summarize each worker's messages in the two categories separately and calculate the difference between those summaries. We then average these differences across workers, giving each worker equal weight. Bars show 95% confidence intervals.

These prompt patterns suggest that workers could be borrowing or relying on expertise from AI to do the cross-occupation work that they are attempting. Fewer requests for explanations or how-to guidance could mean they want to use specialized knowledge without necessarily learning how to do the task themselves. Their greater use of examples and background information may reflect efforts to apply knowledge from colleagues, reference documents, or earlier AI conversations. More checking and verification could signal a need for a second opinion on tasks where workers feel less confident.

See Methodology and definitions for sample construction, occupation mapping, and task classification.

The rising share of cross-occupation tasks

When workers start experimenting with cross-occupation tasks, there are potentially three different long-term outcomes:

- In the first scenario, workers could find that such experimentation is either unproductive or unwelcome within their organization. In that case, their commitment to performing cross-occupation tasks would likely begin to recede in subsequent months.
- In the second scenario, cross-occupation tasks stabilize at a steady percentage of a worker's overall activity, potentially reflecting a normalizing balance between within-occupation and cross-occupation work.
- In the third scenario, workers' initial explorations involving cross-occupation tasks could prove rewarding enough that such activity becomes an increasing part of their AI use over time.

Figure 3 examines whether previously used cross-occupation tasks take up a larger share of a worker's AI activity, finding that cross-occupation tasks become recurring parts of their AI use. To examine this pattern, we use repeated cross-occupation activity among about 6,200 workers observed during our four-month study. We count a message as repeat use when we have previously observed the same worker using that cross-occupation task category. For each worker, we calculate these messages as a share of their sampled non-generic messages.[2] Averaging across workers, this share increased by almost 2x, from 13.1% in April to 25.9% in July.[3]

Figure 3

Previously used cross-occupation tasks become a larger share of occupation-specific work

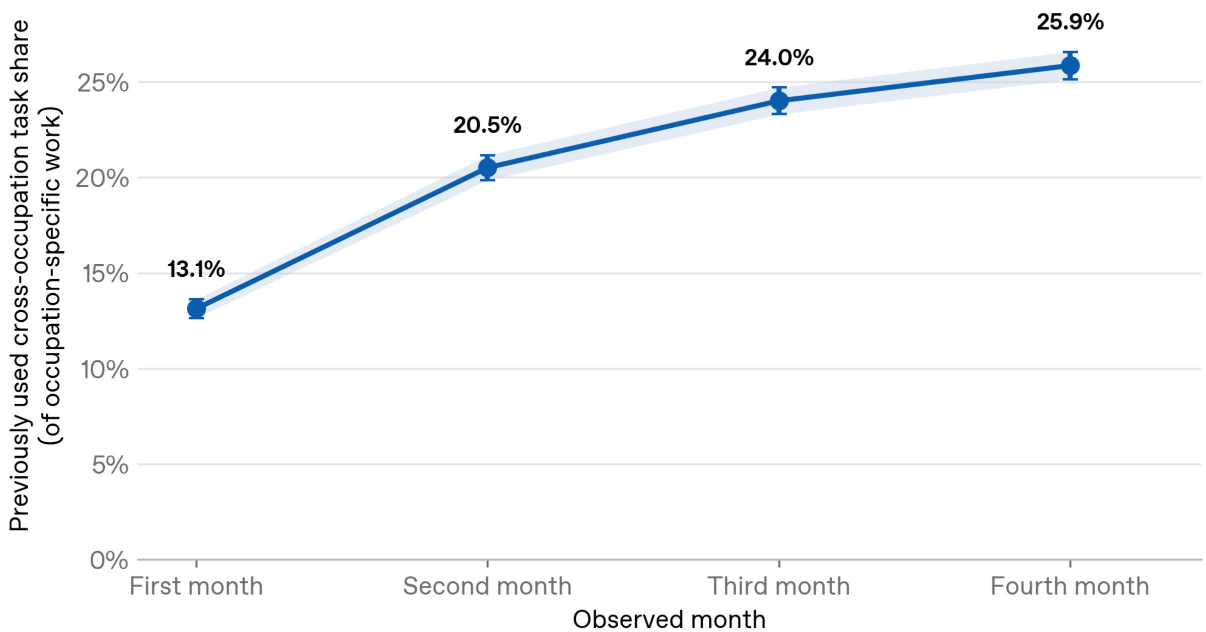

Note: This chart uses the full consistently observed cohort. Shading and bars show 95% confidence intervals.

This pattern is consistent with workers incorporating cross-occupation assistance into their ongoing workflows, potentially broadening the activities they undertake as part of their jobs.

[2] Non-generic messages refer to messages that are classified to an occupation. These exclude messages such as writing, summarizing, and scheduling that are shared broadly across occupations.

[3] As workers try new tasks in any given month, the number of tasks that can be considered recurring increases.

Task recurrence and subsequent coworker use

To interpret these patterns, we need to know how often comparable workers use the same task. We next examine how these task categories reappear in later months, both for the same worker and for coworkers.

Figure 4 asks whether a worker is more likely to use AI for a task in a given month if they—or a coworker— performed the task using AI in the previous month. First, we compare workers who used AI for the task with similar workers who did not. Second, we compare workers whose coworkers used AI for the task with similar workers in workplaces where no one did. In both comparisons, we measure use of that same task in the following month.

Figure 4

Task recurrence and coworker use relative to matched benchmarks

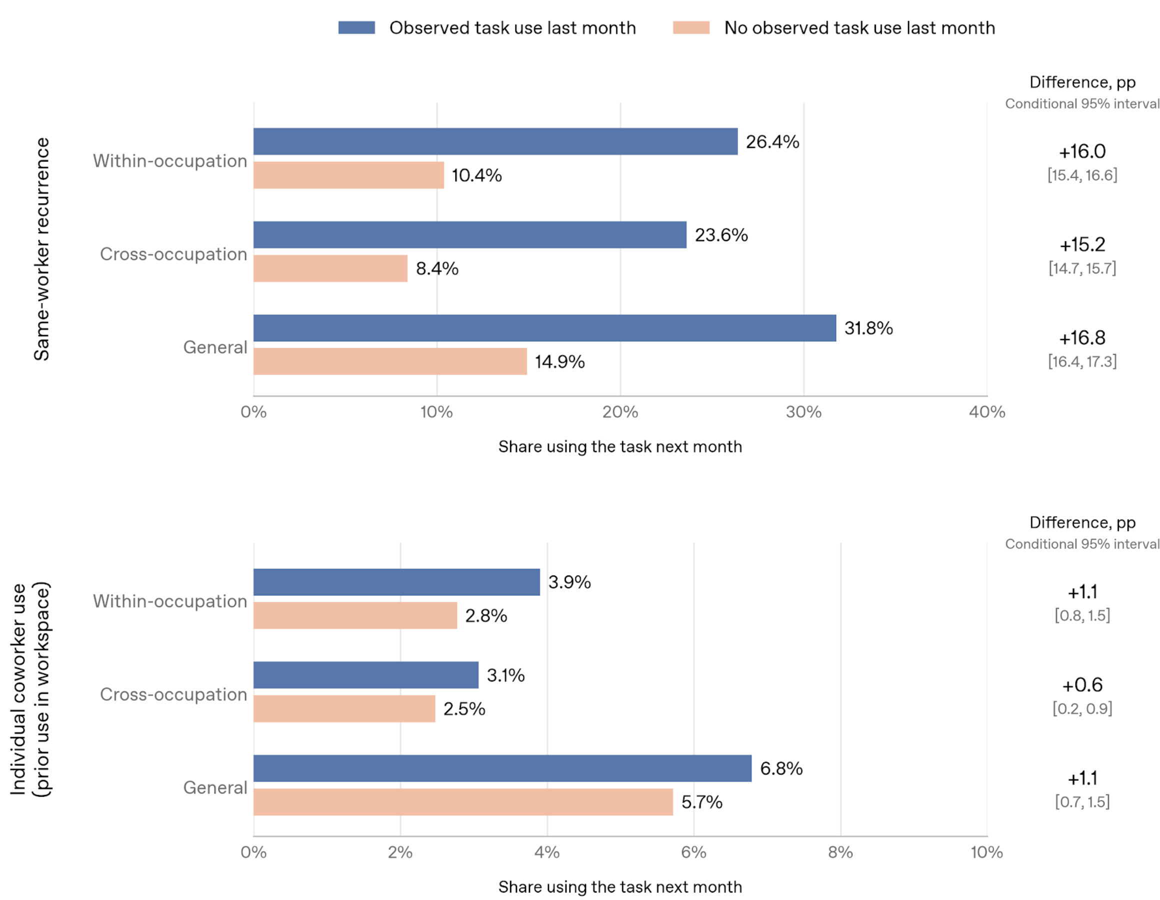

To put recurrence in context, we examine whether workers observed performing a cross-occupation task return to it the following month. We compare each worker's use of the task in the succeeding month with that of workers who had no observed use of the task in the initial month. We match workers in the same broad occupation who have similar levels of AI activity in the following month.

Workers observed using AI for a cross-occupation task in one month were 15.2 percentage points more likely to be observed using it the next month than similar workers with no previous recorded use of that task (23.6% versus 8.4%). The pattern was similar for within-occupation and generic tasks, with gaps of roughly 15–17 percentage points across all three categories. These patterns are consistent with cross-occupation assistance meeting recurring work needs.

We also examine whether these activities subsequently appear in coworkers' AI use, finding that when one person in a workplace does a task, their peers are more likely to as well. When a cross-occupation task appears in a workspace, coworkers without observed use of it that month use that same cross-occupation task the next month in 3.1% of matched cases, compared with 2.5% among comparable workers in other workspaces with no observed use in the starting month.

Together, these findings show strong recurrence within workers and a smaller increase in subsequent use among coworkers. Paired with the previous result, these findings suggest that AI-assisted work across occupational boundaries can become part of a worker's regular routines. The smaller increase in a coworker's subsequent use suggests that the workplace may also shape which activities take hold, whether through shared work needs or colleagues learning from one another.

For organizations, this creates opportunities to rethink how tasks are divided across employees, and where specialist support remains most valuable. The fact that coworkers are also more likely to undertake these tasks could mean that this individual generalization of work accompanies a broader expansion of what teams undertake together.

Cross-occupation tasks that become part of regular workflows

Figure 5 plots the recurrence rates of three distinct tasks that occur most and least often in prompts among widely used cross-occupation tasks such as marketing and promotional work, explaining financial information, and researching legal issues. These differences help identify which activities become repeated parts of workers' AI use and which remain more episodic.

For each task, we check whether each worker returns to the same cross-occupation task in the month after their first observed use of that task. Across all eligible worker–task combinations, a return occurs in 18.5% of cases—roughly one in five. The dotted line marks this overall rate.

Figure 5

Some cross-occupation tasks recur much more often than others

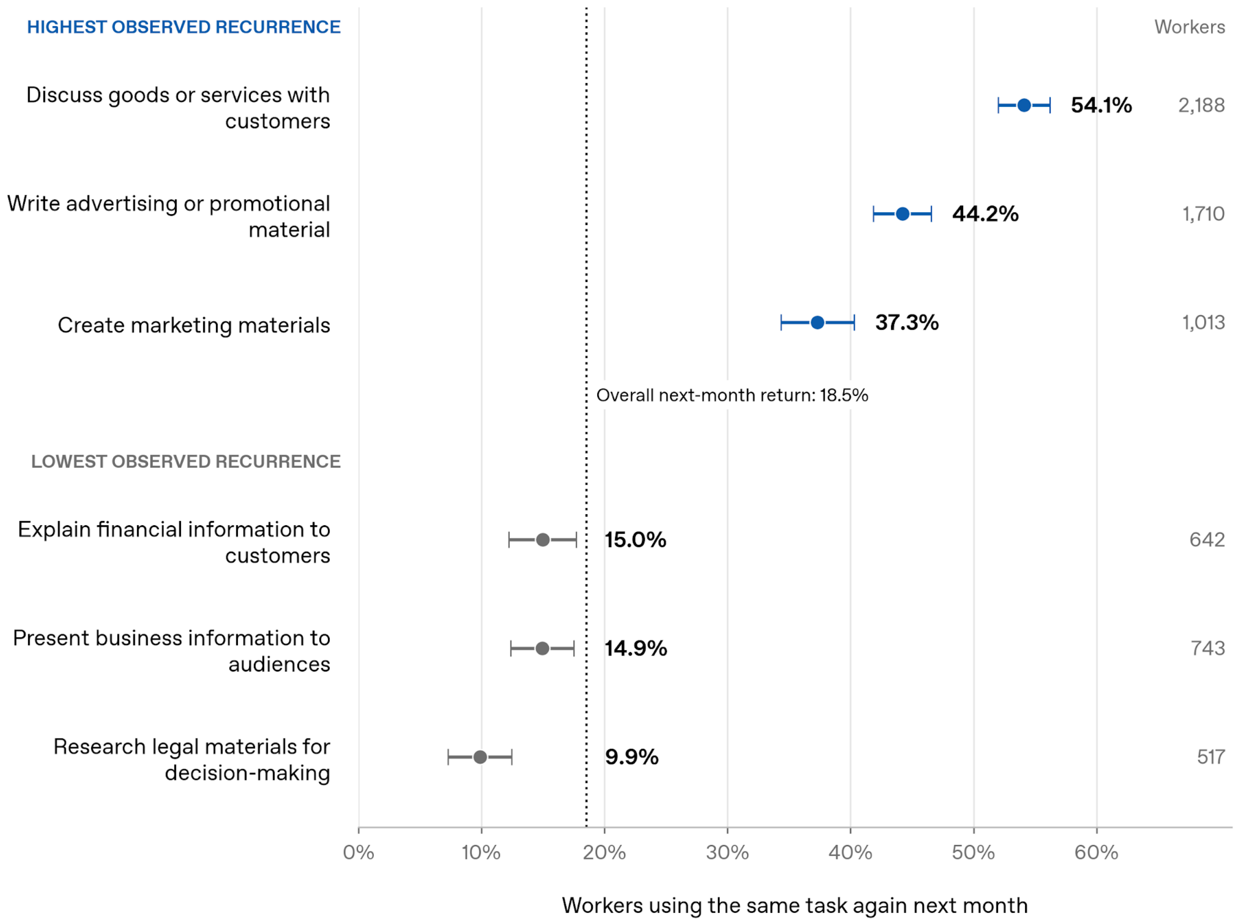

Figure 5 measures whether workers return to a cross-occupation task in the month after their first observed use in our sample. Each worker is counted once per task. Across all eligible worker–task combinations, 18.5% return the following month; the dotted line marks this overall rate. Figure 4's 23.6% measures return after any month of observed use, including subsequent uses, so the same worker–task combination can contribute more than once. Figure 4 also restricts its estimate to cases with a matched comparison.

Our research finds that customer communication (54%), advertising or promotional writing (44%), and creating marketing materials (37%) have the highest observed return rates. Explaining financial information to customers (15%), presenting business information (15%), and legal research (10%) have the lowest observed return rates.

This pattern is consistent with uneven expansion of a typical worker's task repertoire: some activities recur, while others remain occasional. This suggests that AI may redraw some occupational boundaries more in some areas than in others, creating new combinations of tasks within jobs. The resulting changes in specialization and coordination could, therefore, vary substantially across occupations, depending on which outside activities workers repeatedly take on.

What these patterns tell us

AI lets workers “borrow expertise” by drawing on assistance associated with another field. Compared with the same workers’ within-occupation requests, cross-occupation prompts are less likely to seek explanations or advice, and are more likely to provide examples or background and ask for verification. Workers may be bringing the problem and its context, then asking AI for help applying knowledge from another field.

For many workers in our study, borrowing expertise became a more regular part of their work. Workers return to previously used cross-occupation tasks more often than comparable workers with no recorded use of the task in the prior month. This is consistent with some cross-occupation assistance meeting recurring work needs.

Such strategic borrowing offers a way to think about job transformation: changes in the mix of activities *within* a job. If recurring cross-occupation activities become part of a worker’s regular responsibilities, a job could broaden even while its title stays the same. For workers and organizations, the question is therefore not only whether AI can help with tasks beyond an occupational boundary, but which uses become routine and where specialist judgment remains necessary.

For companies that are endeavoring to determine how best to adopt and make use of AI in their organizations, the trends identified in our four-month study of professional worker AI use and prompts represent an important window into how these companies might reduce the friction between identifying a problem and moving their overall work forward aided by AI. Our findings suggest that work design deserves a place alongside access to AI tools in how organizations implement their AI strategies.

Methodology and definitions

We analyze more than 1.5 million work-related messages sent on individual ChatGPT accounts by U.S.-registered users from April through July 2026. We identify users’ occupations and workspace memberships using role or department information and account records linked from ChatGPT Business. We sample conversations at random and retain messages classified as work-related, excluding users who have opted out of training and messages marked as training-disabled. For analyses of changes in workers’ task mix and their return to previously used tasks, we further restrict the sample to workers with at least 10 sampled messages in each of the four months. This activity requirement is applied after sampling and does not mean that we observe all of those workers’ messages. Recurrence is measured within the sampled messages, so repeated task use in unsampled conversations may be missed. The findings describe observed ChatGPT use in this selected population and are not representative of the U.S. workforce.

We analyze work-related messages from individual ChatGPT accounts belonging to U.S. users with parallel ChatGPT Business accounts. We inferred users' occupation and membership in a given workspace from information they reported during onboarding onto ChatGPT Business. We then linked that information to messages from individual ChatGPT accounts. To give an example, assume that user A has both an individual ChatGPT account and a parallel ChatGPT business account. When onboarding onto the latter, user A provided information about their occupation and workspace. We then link that information to user A's individual ChatGPT account and the messages shared on that account.

We include only users whose training-permission records allow their data to be used for model training. Users who have opted out of training are excluded from the linked sample, and we additionally exclude messages explicitly marked as training-disabled. No messages that were sent through a ChatGPT Business account are included in the sample.

Before classification, all message text undergoes automated scrubbing to remove personally identifiable information (PII). Automated models assign task and interaction labels, and the analytical dataset uses hashed identifiers to measure patterns across messages, users, and months without including conversation text. Researchers analyze these labels and derived measures. The research team does not read individual conversations, whether raw messages or PII-scrubbed message text, as part of this study.

We use O*NET, the U.S. Department of Labor's database of occupations and work activities, to identify the tasks traditionally associated with each occupation group. We classify work-related messages to the single detailed work activity (DWA) that most closely matches their task content, where a suitable activity can be identified. We then compare the occupations associated with that activity with the sender's stated occupation, distinguishing within-occupation, cross-occupation, and generic work, with other work left unmapped. The within-occupation/cross-occupation comparisons cover nine occupation groups. This is a sample of ChatGPT users with linked occupation information, not a representative sample of the U.S. workforce.

An occupational boundary describes the activities historically linked to an occupation. Within-occupation and cross-occupation refer to tasks within or beyond that boundary. A cross-occupation task may already be part of a worker's job. A task is an activity category, not an identical assignment. All activity measures refer to observed AI use, not total work or time spent working. Sample and observation counts are rounded to about two significant figures; percentages and confidence intervals use unrounded data.

Appendix: counting all observed task activity

This check uses repeated cross-occupation messages from the highest-activity third of Figure 3's eligible cohort, but divides by all observed task activity—including broadly shared and unmapped tasks.

Figure A1

The pattern remains when all observed tasks are counted

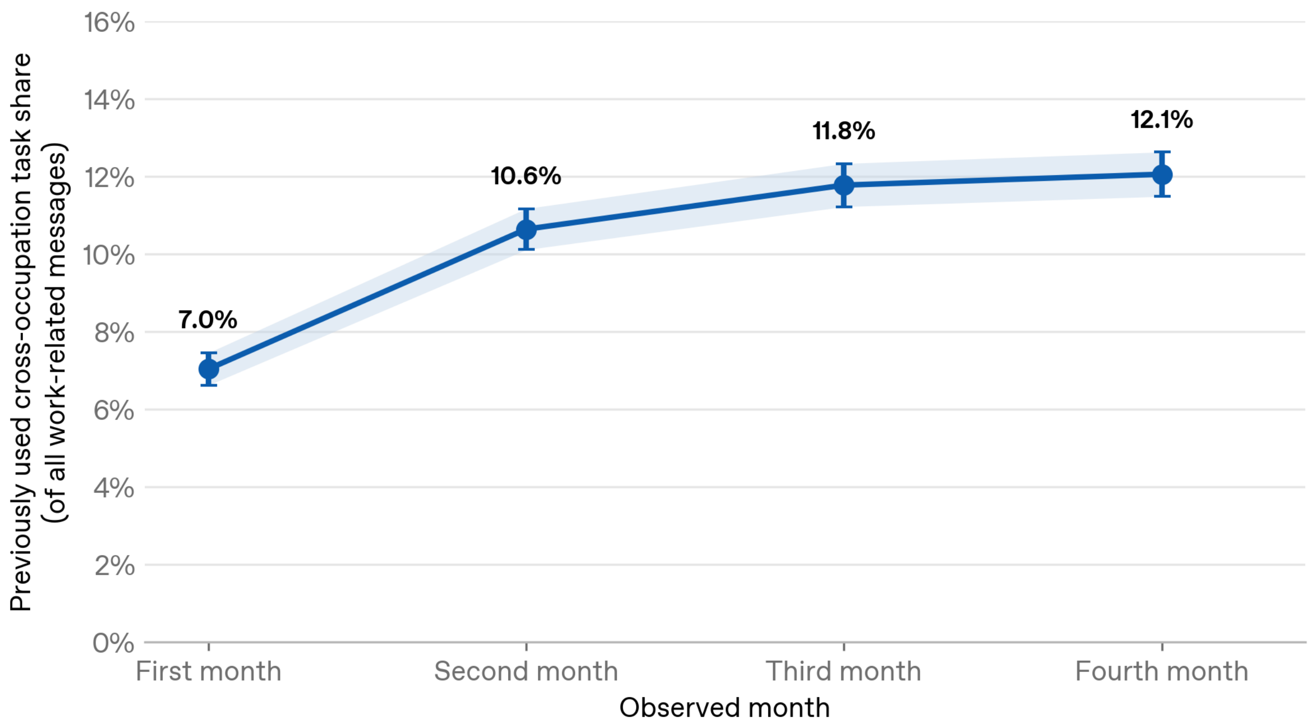

This uses the 2,100-worker highest-activity subset of Figure 3's eligible cohort and equal worker weights, with a broader denominator. The shaded band and bars are 95% confidence intervals.

The average share of previously used cross-occupation tasks rises from 7.0% in April to 10.6% in May, 11.8% in June, and 12.1% in July. Counting all tasks lowers the level but leaves the upward pattern. The same growing-history and selected-sample qualifications apply.

Months 1–4 in the portfolio figures are April–July 2026. All activity measures refer to observed AI use, not total work or time spent working.

OpenAI

Thank you

Keep exploring work at the frontier

OpenAI Signals

Explore data and research on AI's impact

openai.com/signals

The Economics of AI

Get AI economics insights in your inbox

economicsofai.substack.com