

Work at the Frontier: How AI is expanding what people do at work



Key takeaways

- 01** AI may change not only how work gets done, but who does it. Our new research suggests that ChatGPT users frequently seek help with tasks historically associated with other occupations, suggesting AI may broaden roles and change how work is divided.

- 02** Across our sample, 16.8 percent of all work-related messages—and 43.5 percent of occupation-specific messages—concern tasks historically associated with another occupation.

- 03** Patterns of AI use differ across occupations. Marketing and engineering tasks travel to many workers in other occupations, while workers in design, sales, and human resources are more likely to take on a variety of tasks from other jobs.

- 04** These effects are more pronounced at smaller firms. Typical users at smaller workspaces do relatively more cross-occupation work. This could be a sign that users turn to AI as a tool to expand the types of work they do when they have fewer resources.

Introduction

How is AI changing how we work? Our analysis of work-related messages sent by ChatGPT users provides a window into the changing task composition of jobs. We find that 43.5 percent of work-related, occupation-specific messages concern tasks historically associated with another occupation.¹

In our *AI Jobs Transition Framework*, we highlight that a significant portion of jobs are likely to reorganize, that is, their day-to-day tasks may shift considerably. Our *Work at the Frontier* series is aimed at understanding this process. As more and more workers use AI in their jobs, new combinations of individual and machine activity that weren't previously possible emerge. Understanding these changes is at the heart of understanding how work adapts to AI: it changes what a worker does alone, what they delegate to a colleague, and what requires a specialist.

With AI, a salesperson who once handed a customer dataset to an analyst can now easily and quickly explore it themselves. A marketer who once waited for a software or web developer can now troubleshoot a website or write a simple script. In each case, AI changes both how quickly work is completed and who does it. We examine how AI expands the types of work people do, especially when it is not within their typical job description. We find that AI may change the historical or typical boundaries of what workers do within a certain role.

An occupational boundary is the set of activities historically associated with an occupation. We examine whether workers use AI for activities within or outside that boundary.


We focus on users in eight occupation groups: customer experience, design, engineering, finance, human resources, legal, marketing, and sales. 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 use a random sample of work-related messages from these users and classify each message to a single O*NET detailed work activity (DWA) that most closely matches its task content. We then compare the occupations associated with that activity with the sender's stated occupation.² Occupational boundaries are not always precise because generic tasks like writing, summarizing, scheduling, and many other activities appear across many types of jobs. We therefore assign messages to one of three categories: Generic, Within occupation, and Cross-occupation.

¹ We link users to occupation using role information provided by the users' ChatGPT Business accounts.

² We drop users from the sample who do not provide a Department and Role when they sign up for ChatGPT Business. We use ChatGPT to anonymously classify each message according to its task content and never read underlying messages.

Fig. 1

Identifying task crossover



Note: Occupational boundaries are based on tasks historically associated with each occupation in O*NET.


In this random sample of work-related messages:

- 16.8 percent of messages are classified as Cross-occupation, meaning they concern tasks historically associated with another occupation;
- 21.8 percent are classified as Within occupation, meaning they concern tasks historically associated with the user’s own occupation;
- and 61.5 percent are classified as Generic, meaning they relate to tasks that are broadly shared across occupations like writing emails or scheduling meetings.

Among non-generic messages, 43.5 percent are cross-occupation. Excluding generic work like writing emails, this category accounts for 28 to 77 percent of occupation-specific messages.

Fig. 2

Nearly half of occupation-specific AI use is cross-occupation



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET DWA descriptions. Notes: Occupation-specific shares exclude generic work and rebase the remainder to 100%.

We call this pattern *task crossover*: work historically associated with one occupation appearing in the AI use of people in another. Across the eight occupation groups, between 11 and 30 percent of messages involve task crossover.

Task crossover differs between different types of workers. Marketing and engineering work appears widely in the AI use of workers in other fields. Meanwhile, workers in fields like design, sales, and human resources are more likely to use AI for tasks outside of their typical roles. Task crossover is more common in smaller workplaces, but the difference is concentrated among typical-volume users: cross-occupation use is 18.9 percent in 2–5 seat workspaces and 16.3 percent in workspaces with 101 or more seats among users in the middle 50 percent by message volume; the analogous pattern is not monotonic among top-quartile users. This split could reflect the possibility that workers seek AI help with a broader range of tasks where fewer colleagues or specialists are immediately available.

Our results point to a broader conception of AI-driven occupational change. AI may automate some existing tasks and augment workers in others. AI may also change how work is divided by broadening the set of tasks workers can do and allowing work associated with one occupation to migrate into another. For labor-market policy, workforce development, and the organization of firms, it is critical to understand which new bundles of human and machine activity become viable, and who is prepared, permitted, and paid to perform them.

Background

In labor economics, the task-based framework shifted attention from occupations as indivisible units toward the activities that workers and machines perform. Autor, Levy, and Murnane (2003) use a task-based framework – examining the individual tasks that make up users’ jobs – to understand how different occupations were affected by the advent of computers. Later, Acemoglu and Restrepo (2019) emphasized that technology changes demand for workers in different roles by reallocating existing tasks to capital and creating new tasks in which labor has a comparative advantage. This approach made clear that technology need not affect every part of a job in the same way and produced a framework that can be used to analyze how different factors, like technology, affect jobs by changing tasks.

Accordingly, most applications of the task-based framework to AI begin with the bundle of activities currently assigned to an occupation and ask how those activities will change with technology. However, this assumes that the set of tasks each worker does is fixed. Gans (2026) discusses this limitation. When automation changes the cost of performing particular activities, firms may split an existing job into narrower roles, allowing workers to specialize, or even recombine activities into a broader role, creating more generalists. The key point is that the bundle of human work – and therefore the bundle of skills rewarded in the labor market – changes in response to AI.

Historical evidence suggests that occupation boundaries are continually evolving. Atalay, Phongthientham, Sotelo, and Tannenbaum (2020) find that a substantial share of long-run change in the task content of work occurred within job titles, not just through changing occupation composition. Autor, Chin, Salomons, and Seegmiller (2024) show that technological change also can produce new forms of work that eventually appear in occupational classifications. Those studies observe changes to jobs and tasks after this has been incorporated into job advertisements, titles, or employment patterns.

Analysis of Generative AI usage offers a window into an earlier stage of the job evolution process, as workers can begin testing and recombining tasks *before* firms revise job descriptions or invent new job titles altogether. Recent AI-use evidence points to this shift: Yang et al. (2026) find that queries on Perplexity regularly fall outside of users’ inferred primary occupation, suggesting that AI tools may expand the scope of work users attempt. We measure task crossover more directly using work-related messages linked to self-reported occupations, allowing us to see when users seek AI help with tasks traditionally associated with another occupation and how these patterns vary with workplace size.

Workers use AI for tasks beyond their traditional occupations

Workers frequently use AI for tasks beyond their traditional occupational boundary. Workers ask AI to do work like troubleshooting software, drafting advertising copy, and explaining technical specifications, even when those activities are not central to their own occupation.

With each message classified as Within occupation, Generic, or Cross-occupation, we examine how task crossover varies across the eight user occupations.

Table 1


Task crossover classification

Within occupation	Directly linked to, or near, the user’s occupational boundary.
Generic	Common across many occupations.
Cross-occupation	Linked to another occupation; counted as boundary crossing.

Cross-occupation tasks make up most occupation-specific messages in five of eight groups: 77 percent for customer experience, 75 percent for design, 69 percent for human resources, 56 percent for legal, and 53 percent for marketing. Figure 3’s denominator is occupation-specific (non-generic) messages within each worker occupation: it excludes broadly shared work and rebases the remaining within-occupation and cross-occupation shares to 100 percent.

Fig. 3

Cross-occupation work is a majority in five of eight groups



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET DWA descriptions. Notes: Each message is assigned to the occupation whose traditional tasks it most closely resembles. The denominator is occupation-specific (non-generic) messages within each user occupation; within-occupation and cross-occupation shares are rebased to 100%. Bars show only the cross-occupation share.

Customer experience, design, and human resources show the highest cross-occupation shares. Next, we turn to understanding which tasks travel across occupations.

Workers borrow specific tasks from other jobs

Some tasks appear repeatedly in messages across users in many different occupations. In our random sample of work-related messages, recurring tasks include creating marketing materials, troubleshooting computer applications, calculating financial data, and explaining regulations or policies.

Financial calculation and computer troubleshooting are particularly common tasks across user occupations. Calculating financial data ranks among the three most common finance-related tasks in every non-finance user occupation. Similarly, troubleshooting computer applications and systems ranks among the top three engineering related tasks in every non-engineering user occupation. Of all finance-related messages sent by non-finance users, 9.8 percent concern calculating financial data. Of the engineering-related messages sent by non-engineering users, 6.1 percent concern troubleshooting computer applications or systems.

Among users in non-marketing occupations, developing promotional materials is the most common marketing-related task, accounting for 25 percent of marketing-related messages. AI-generated summaries of messages assigned to this category include requests to create ads, social posts, flyers, promotional videos, presentations, product sheets, and graphics for marketing campaigns, showing the range of work assigned to this task. Creating marketing materials is the next most common marketing-related task, accounting for 9.3 percent of marketing-related messages, with its highest share among designers. Developing marketing plans or strategies accounts for another 6.5 percent.

Table 2

Recurring cross-occupation tasks

Task	Closest occupation	Groups	Top group
Calculate financial data	Finance	7/7	Sales: 14.3%
Troubleshoot issues with computer applications or systems	Engineering	7/7	Customer experience: 7.6%
Discuss goods or services information with customers or patrons	Customer experience	7/7	Marketing: 70.8%
Create marketing materials	Marketing	5/7	Design: 13.6%
Communicate with government agencies	Legal	7/7	Customer experience: 21.4%

Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET descriptions. Notes: “Closest occupation” is the occupation whose traditional work the task most closely resembles. The remaining columns refer to users who belong in the seven other occupation groups. “Groups” counts the number of the remaining seven user occupations in which the task ranks among the three most common detailed work activities related to the closest occupation. Top group is the user occupation in which the listed task makes up the largest percentage of messages matched to the closest occupation.

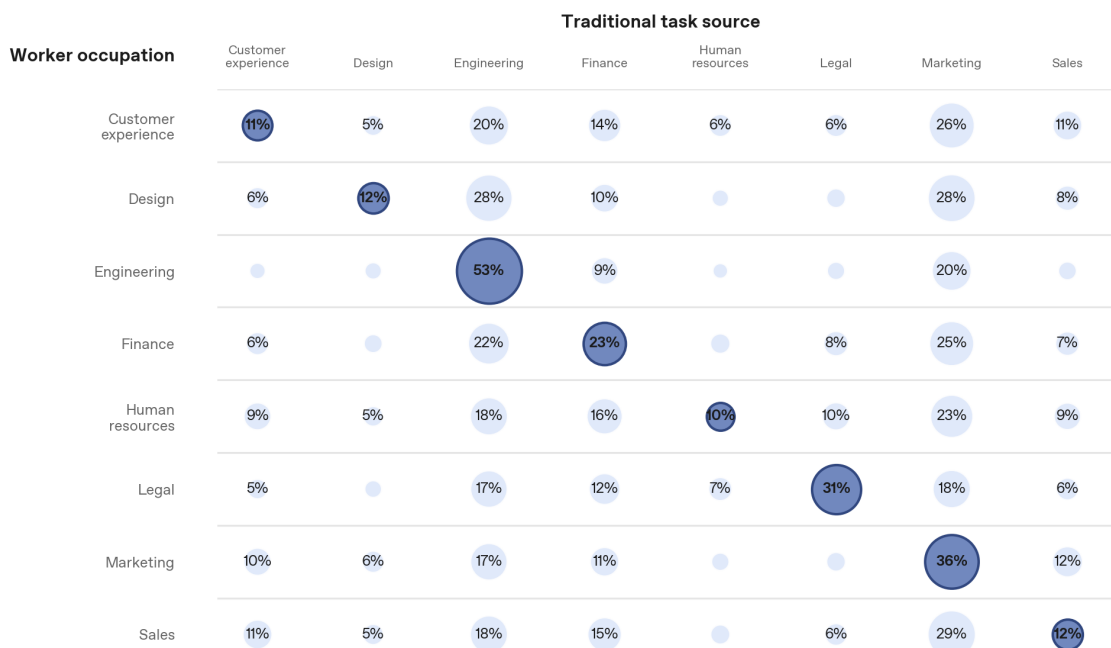
Marketing and engineering tasks travel across jobs

Figures 2 and 3 show whether a non-generic activity falls within the user's occupational boundary or outside it. Because occupational boundaries overlap, an activity may fall within more than one occupation's boundary. Figure 4 assigns each activity to a single occupation whose tasks it most closely resembles. For each user occupation, the corresponding row reports the share of non-generic messages whose task content most closely matches each column occupation. Workers in fields other than marketing and engineering do work associated with those occupations most frequently. Marketing tasks account for about 28 to 29 percent of non-generic messages among sales and design users and 26 percent among customer-experience users. Engineering tasks account for 28 percent among design users and about 20 to 22 percent among customer-experience and finance users.

Bubble area shows the share of a worker group's occupation-specific messages that are linked to each traditional task source. Dark diagonal bubbles show each source's own-occupation benchmark. Displayed row percentages may not sum exactly to 100% due to rounding.

Fig. 4

Marketing and engineering tasks travel across many occupations




Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET descriptions. Notes: Messages are assigned a traditional task source. The denominator is each worker occupation's non-generic, occupation-specific messages; Generic work is excluded.

Some occupations draw on cross-occupation tasks; others have tasks that travel

Figure 5 separates the two directions of task crossover, unpacking which occupations send more tasks to other occupations and which receive them. The left panel uses all work-related messages within each worker occupation as its denominator, including generic messages; the right panel shows how often each occupation’s tasks appear in messages from users in other occupations:

- Tasks brought in: the left panel shows how much of each worker group’s AI use draws on other occupations
- Tasks that travel: the right panel shows how widely each occupation’s own task bundle appears elsewhere.

Fig. 5
Design draws most on cross-occupation tasks; marketing and engineering travel furthest



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET DWA descriptions. Notes: The left panel uses all work-related messages within each user occupation as its denominator, including generic messages, and counts the share assigned to other occupations’ traditional tasks. The right panel shows how often each occupation’s tasks appear in messages from users in other occupations. For each occupation, we calculate the share of messages related to it among users in each of the seven other occupations. We then report the average of those seven shares.

The two kinds of task crossover do not always go together. Workers in an occupation may often use AI for tasks normally associated with other fields, even when workers in other fields rarely take on that occupation's tasks. The reverse can also be true.

Design illustrates the first pattern. About 35.2 percent of messages sent by designers involve work usually associated with another occupation. But design tasks make up only 1.7 percent of the messages sent by workers in other fields. In other words, designers draw heavily on tasks from elsewhere, while design work itself rarely appears outside design.

Engineering is closer to the opposite pattern. Only 18.5 percent of engineering messages involve tasks from other fields, but engineering tasks account for 7.4 percent of messages among workers in other occupations. Engineering therefore appears to be an important source of tasks that workers elsewhere adopt.

Marketing stands out in both directions. Marketers devote 24.3 percent of their messages to tasks associated with other occupations. At the same time, marketing tasks account for 8.9 percent of messages among workers in other fields—the highest share in the sample.

Figures 4 and 5 therefore reveal three different roles in the movement of tasks across occupations. Designers use AI to combine work from several fields. Engineering tasks are frequently taken up by workers elsewhere. Marketing does both: marketers take on a substantial amount of cross-occupation work, and marketing tasks also appear widely across other occupations.


These patterns suggest that AI may change how work is divided across occupations, but this does not mean that the occupations themselves are disappearing. AI may make a task easier for an outsider to attempt while specialists remain critical for expert-level judgment and review. These results should therefore be understood as evidence that the division of work is shifting, not as direct estimates of which occupations will gain or lose jobs.

Smaller workspaces show more boundary crossing among typical-volume users

Figure 6 focuses on typical-volume users: users in the middle 50 percent of message volume. Its denominator is all work-related messages among typical-volume users, including generic messages. For this group, the share of messages classified as cross-occupation falls from 18.9 percent for users in 2–5 seat workspaces to 16.3 percent for users in workspaces with 101 or more seats. This is about 2.5 percentage points, or roughly 13 percent relative to the rate in the smallest workspaces. Among high-volume users, by contrast, there is no monotonic trend.

Fig. 6

Smaller workspaces show more boundary crossing among typical-volume users



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET descriptions. Notes: The denominator is all work-related messages within each workspace-seat bucket, including generic messages. Cross-occupation messages are beyond the user's O*NET-derived boundary and are not classified as generic. Shares are message-weighted, so higher-volume users contribute more observations. The chart shows users in the middle 50 percent of the user message-volume distribution. Workspace seats are not total company size.

The greater boundary crossing observed in small workspaces appears among the more typical users who make up the middle 50 percent of the usage distribution. One possible interpretation is that occasional or moderate users in small organizations turn to AI when they encounter tasks that would otherwise require help from another worker. A worker in a small business may use AI intermittently to draft marketing copy, troubleshoot software, review a contract, or perform basic analysis because there is no specialist to delegate to. In larger organizations, the same worker may be more likely to rely on an established team, workflow, or internal service.

Repeating this analysis among heavy users – users in the top quartile by message volume – reveals a different pattern of use. Heavy users may have developed stable AI-supported workflows that are relatively similar across organizational settings. Alternatively, their additional message volume may reflect repeated work within their core occupation—such as iterative coding, editing, or analysis—rather than continued expansion into new functional areas. In that case, higher message volume would increase the amount of AI-assisted work without necessarily increasing the cross-occupation share.

These comparisons are descriptive. Workspace size may be a proxy for many factors – like specialization, industry, user occupation, organizational maturity, and access to colleagues. This also is distinct from total company size, as firms may purchase a different number of seats than they have total employees.

Conclusion

Many studies of AI and work begin with a fixed list of tasks and ask which ones a model can perform. Our evidence suggests that the list itself is changing. People bring AI into work that lies outside the historical boundaries of their occupation. An open question is whether repeated AI use for activities outside workers' traditional occupational boundaries will lead to lasting changes in job responsibilities. If it does, workers may need training to evaluate AI-assisted work outside their established expertise, and organizations will need to prepare clear processes for review and accountability.

This report also highlights a measurement challenge. Government datasets provide a valuable account of how work has historically been divided across jobs. If AI changes which workers perform particular tasks, measures based only on existing job descriptions will gradually diverge from how work is actually organized.

Methodology and limitations

We analyze a random sample of over 800,000 work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations and account information were linked from ChatGPT Business; we filter these messages to those we classify as work-related. The role data come from ChatGPT Business; the analyzed messages come from those users' individual ChatGPT accounts. ChatGPT Business and Enterprise are separate product populations, so these estimates should not be generalized to Enterprise users. The analysis covers eight occupation groups. We compare each user's stated occupation with the main work activity in each message using O*NET as a historical baseline for the activities associated with different occupations.

This study maps AI-aided work; it does not estimate AI's effect on employment or productivity. The random sample is not representative of the entire U.S. workforce. Occupations come from self-reported ChatGPT Business role information; the analyzed messages are work-related activity on those users' individual ChatGPT accounts.

The unit of analysis is a message, not an hour of work, a completed project, or a job. We do not observe whether the output was used, how good it was, how much time it saved, whether the user could have completed the task without AI, or whether a specialist reviewed it.

Eligible messages are first assigned one intermediate work activity (IWA). They are then mapped to a detailed work activity (DWA) drawn from that IWA or from one of the five most semantically similar IWAs, based on IWA title embeddings. Occupational boundaries are then constructed from an employment-weighted occupation-to-SOC crosswalk, O*NET task-to-DWA links, task ratings, and DWA embeddings. Because a message may contain multiple or overlapping tasks that could correspond to different DWAs and occupations, the assigned DWA represents its primary activity rather than every task or occupation reflected in the message.

Most importantly, these results are descriptive. They show that people use AI for a wider range of activities than a fixed job description would predict, but they do not reveal how the same people would have allocated their work without AI.

Technical appendix: classification and measurement

A1. Sample and unit of analysis

The initial sample contains more than 800,000 work-related messages sent from the individual ChatGPT accounts of U.S. users. We assign each user to one of eight occupation groups using self-reported Department or Role information from ChatGPT Business and retain only users whose information maps to a single group. The unit of analysis is a user message. To classify each selected message, the model uses that message and up to nine preceding messages from the same conversation as context; only the selected message is classified.

A2. Hierarchical classification and occupational boundaries

Classification proceeds in three steps. First, a model determines whether a message is eligible for occupational-task classification, which means that we can identify the occupation of the user, the user is based in the U.S., and the message is classified as work-related. Eligible messages are first classified to one IWA (Intermediate Work Activity). They are then classified to a DWA drawn from that IWA or from one of the five most semantically similar IWAs, based on IWA title embeddings. This is hierarchical rather than a single direct message-to-DWA label.

For each user occupation, we construct an O*NET-based task boundary. We map the user occupation to detailed SOC occupations using employment-share weights, then use O*NET task links and task ratings to identify the DWAs associated with those SOCs. Each DWA is embedded using its DWA title plus its parent IWA title. A DWA falls within an occupational boundary if it is directly linked to the occupation or when its cosine similarity to the closest boundary DWA is at least 0.80.

A3. Reader-facing categories and task sources

After each message is mapped to a detailed work activity, we assign it to one of three categories. First, we label a message “Generic” when its detailed work activity resembles work traditionally linked to a majority of the occupation groups used in classification. We also make a limited set of manual adjustments by classifying broadly shared activities, such as editing written materials and gathering information about goods and services, as generic; reassigning activities only when manual review identifies a clearly more suitable occupational match; and marking activities as outside the occupation taxonomy when they do not fit any occupation used for matching. For the remaining messages, a message is labeled “Within occupation” if its task is directly linked to, or semantically close to, the user’s own occupational boundary. All other remaining messages are labeled “Cross-occupation.”

The common analysis sample used to construct the figures includes generic messages and non-generic messages assigned to one of the eight occupations studied. Non-generic messages that fall outside those eight occupations, for example messages assigned to healthcare, are excluded. Each figure then applies the additional restrictions described below.

For figures that show the occupation traditionally associated with a task, we assign each non-generic DWA to one displayed occupation using the same O*NET-based boundaries described above. If the DWA is directly linked to one displayed occupation, we assign it to that occupation. If it is directly linked to more than one displayed occupation, we assign it to the occupation where that DWA has the largest boundary weight. If there is no direct O*NET link to a displayed occupation, we compare the

DWA with the tasks already associated with each displayed occupation in O*NET, and assign it to the occupation with the closest match when the similarity score is at least 0.80. This measure describes which occupation's O*NET tasks the DWA is closest to, not whether the work was delegated from, or actually originated in, that occupation.


A4. What the figures measure

Figure 3 uses occupation-specific (non-generic) messages within each worker occupation as its denominator and reports the cross-occupation share after rebasing the remaining within-occupation and cross-occupation shares to 100 percent. Figure 4 also uses non-generic, occupation-specific messages within each worker occupation; each cell reports a traditional task source's share of that denominator. Figure 5 shows two directions of task crossover: the left panel uses all work-related messages within each user occupation, including generic messages, while the right panel averages a task source's presence across other user occupations. Figure 6 uses all work-related messages within each workspace-seat bucket, including generic messages, and reports the share classified as cross-occupation for typical-volume users, defined as the middle 50 percent of the user message-volume distribution. Figure A1 shows the same comparison separately for the bottom quartile, middle 50 percent, and top quartile.

Figure A1 shows the full user-volume split behind Figure 6. The middle 50 percent shows a steady small-workspace gradient; the top quartile does not show a monotonic decline, and the bottom quartile falls sharply but is less precise in the largest-workspace bucket.

Fig. A1

Business-size differences vary across user-volume groups



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET descriptions. Notes: Panels split users by sampled message volume: bottom quartile, middle 50 percent, and top quartile. Shares are message-weighted and unadjusted; the largest-workspace cell in the bottom quartile contains 197 messages, so that panel is less precise. Workspace seats are not total company size.

Figure A2 includes generic work in the same-occupation diagonal. Each cell shows the share of messages within a worker occupation assigned to the traditional task source shown in the column; diagonal cells combine same-occupation and generic work.

Fig. A2

Task-source mix with shared work included on the diagonal



Source: OpenAI Economic Research analysis of a random sample of work-related messages on individual ChatGPT accounts from U.S. users whose self-reported occupations were linked from ChatGPT Business account information, with tasks mapped to O*NET descriptions. Notes: Each cell reports a task source’s share within a worker occupation. Diagonal cells combine same-occupation and shared work; percentages are message-weighted.

When we reclassify “generic” work as own-occupation in Figure A2, it makes clear that task crossover does not replace role-specific work: every occupation still retains a meaningful diagonal of same-occupation or shared activity. At the same time, the off-diagonal cells show that users routinely combine their core work with tasks traditionally associated with other fields, especially marketing and engineering. Read this way, the figure is indicative of task bundles becoming more mixed: AI use appears to broaden roles while preserving a recognizable occupational center.