

ANTHROPIC

Anthropic Economic Index report: Cadences

Published

June 26, 2026

Authors

Maxim Massenkoff, Eva Lyubich,
Szymon Sacher, Zoe Hitzig, Shaoyi Zhang,
Ryan Heller, Peter McCrory

Acknowledgements

Scott Booth, Keir Bradwell, Meredith Callan, Dexter Callender III, Boris Cherny, Chris Doenlen, Eleanor Dorfman, Jake Eaton, Evan Frondorf, Deep Ganguli, Romello Goodman, Ankit Gupta, Kunal Handa, Rebecca Hiscott, Andrew Ho, Hanah Ho, Jerry Hong, Saffron Huang, Mo Julapalli, Katie Kennedy, Jennifer Martinez, Miles McCain, Kelsey Nanan, Tyler Neylon, Adnan Pirzada, Dianne Penn, Kerry Persen, Sarah Pollack, Ankur Rathi, Santi Ruiz, David Saunders, Ankit Siva, Michael Stern, Ami Vora, Scott White, Heather Whitney, Kim Withee, Ryan Zauk, Jack Clark.

Introduction

One year ago, most Claude usage took the form of a conversation between a user and an assistant. Today, with the rapid growth of Claude Code and Cowork, Claude sessions increasingly consist of long-running agentic tasks. A chat transcript no longer fully captures how people are using AI, and our methods for studying Claude’s economic impacts have had to rapidly adapt.

To keep pace, we made several changes to our data pipeline for the Economic Index. In this version, we:

- Sample data at a higher rate, allowing us to view usage patterns down to the hourly level.
- Introduce a new classifier that labels the output of each conversation.
- Share more granular data, breaking out results for chat and Cowork conversations (together, “Claude conversations”) and the 1P API, aggregated at a monthly level.¹

We describe additional methodological changes in the Appendix. Together, these changes provide a clearer picture of how AI mirrors and diffuses into economic life.

In addition, we’ve previously lacked visibility into Claude’s impact outside of user sessions. How do people perceive AI to be changing their work, or the opportunities available to them? Does their usage of AI shape their expectations? In an ideal world, what would they want from AI? We report initial findings from the [Anthropic Economic Index Survey](#), launched in April 2026.

We preview our main findings below.

- In Chapter 1, we show how the rhythms of the external world shape Claude usage. Work-related queries subside on the weekend, though less dramatically in the most highly paid occupations; people tend to ask for the news in the morning, and sleep advice peaks around 5 a.m.; tax-related requests surge around filing deadlines.
- Chapter 2 explores the concrete outputs that people take away from their Claude sessions. These are highly dependent on what product they’re using.

Chat and Cowork provide more explanations than Claude Code, for example. The nature of the output also shapes people's interactions with Claude. Building a website leaves much more to Claude's judgment than translating a document, where the answer is largely determined by the text. We also see that more compute is associated with more valuable artifacts; the tokens a given output consumes rise with the estimated value of the work.

- Chapter 3 presents the first results from the Anthropic Economic Index Survey, which we link to Claude usage data through our [privacy-preserving system](#). Expectations and experiences vary systematically with how people use Claude: people who use Claude in the most automated way expect AI to take on more of their tasks in the next year, yet feel the most optimistic about what that means for their work, anticipating positive impacts on pay, job security, and meaning.

¹ This includes chat conversations and Cowork sessions from consumer (Free/Pro/Max) accounts on both [Claude.ai](#) and the Claude desktop app. "First-party API" or 1P API refers to developer traffic routed directly through Anthropic's own programming interface, which is distinct from both Anthropic's consumer-facing Claude.ai application and third-party platforms such as Amazon Bedrock or Google Cloud Vertex. We continue to manage data according to our privacy and retention policies, and our analysis is consistent with our terms, policies, and contractual agreements.

CHAPTER 1

Cadences

Our new privacy-preserving telemetry, which continuously samples a slice of conversations every day, allows us to study daily and hourly patterns in usage, in contrast to the seven-day samples each previous Economic Index report drew on. These analyses capture ebbs and flows in work patterns around the world.¹

We find that Claude usage mirrors the workweek, with personal prompts spiking on the weekend. The hourly data captures within-day patterns—people most often ask for sleep advice around 3 a.m. and for recipes around 6 p.m. We also see usage reflecting key dates. For instance, tax-related requests surged just before the US filing deadline on April 15.

The workweek

The share of chat and Cowork² conversations categorized as personal use spikes from around 35% on weekdays to just under 50% on weekends during the sample period (Figure 1.1). Outside the workweek, users' conversations shift from business correspondence, marketing copy, and slide decks to emotional support, medical questions, and investment advice. This shift is biggest for high-income countries.

The workweek




Figure 1.1: Personal conversations increase on the weekend. Daily share of conversations that are personal use across [Claude.ai](#), Claude Desktop, Claude Code, and 1P API. Saturday and Sunday are shaded grey.

A similar pattern is present in Claude Code and the 1P API traffic (i.e., API traffic routed directly through Anthropic), though both have lower baseline rates of personal use.³

Request clusters⁴ allow us to go one level deeper and see which specific Claude Code tasks swing most between weekdays and weekends. On weekends, the Claude Code usage clusters that fall the most include backend architecture, API debugging, and data storage. Those that increase the most include AI agent design, quant trading, and gaming.

Weekends may also create space for people to pursue new ventures. Across countries, conversations related to starting a business are highest on Saturday and Sunday. However, job application activities drop on the weekend along with other work-related tasks.⁵

Daily rhythms

Hour by hour, Claude usage reflects the rhythms of daily life. Figure 1.2 shows the hourly frequency of different request clusters relative to their overall average in global traffic.⁶

People ask for news at 7 a.m. local time. Business correspondence (e.g., email drafting) traces the arc of the workday, with a slight peak at 10–11 a.m. One of the biggest spikes is recipe requests, which are 2.3 times more frequent at 6 p.m. compared to the average. Media recommendations are most concentrated in the evening, while people seek sleep advice in the few hours just before dawn.




Figure 1.2: Request clusters over the course of the day
Normalized hourly share of conversations that fall into different request clusters, restricted to Claude chat and Cowork data.

On nights and weekends, when people do turn to Claude for work, the tasks skew toward higher-wage occupations (Figure 1.3). While we can't conclusively identify the jobs of the people making these requests, this could reflect the fact that people in higher-paying occupations—like marketing managers or computer programmers—are more likely to work outside traditional hours. In contrast, tasks related to jobs in the bottom two quartiles—like telemarketing and clerical work—fall to a smaller share of total conversations. This pattern isn't driven exclusively by computer and mathematical tasks: when we removed those occupations from the analysis in a robustness check, higher-quartile tasks still increased on nights and weekends.




Figure 1.3: Change in share of work-related conversations during nights and weekends, split by occupation wage quartile. Each bar shows the percent change in the share of work-related tasks coming from the specified wage quartile on nights and weekends versus weekday working hours. Wage quartiles are calculated using BLS data, weighted by number of transcripts.

Tax day

The sample period for this report covers tax filing deadlines for people in the United States. Figure 1.4 shows a large spike in the share of tax-related conversations around the deadline. On April 14, tax-related clusters were eight times as common as on the average day in May and remained about as high on April 15. On April 16, they dropped sharply.

Tax day in Claude traffic




Figure 1.4: Tax-related conversations spike just before the US filing deadline.

Share of conversations related to taxes in the US and the rest of the world.

- ¹ Throughout, all analyses are based on privacy-preserving classifiers: transcripts are only read by another instance of Claude. Then we filter out any cells with insufficient observations to ensure privacy-preserving analysis.
- ² This includes all conversations held in the chat and Cowork tab on both [Claude.ai](#) and through the Claude desktop app. Claude Code and API traffic are presented separately.
- ³ "First-party API" or IP API refers to developer traffic routed directly through Anthropic's own programming interface, which is distinct from both Anthropic's consumer-facing Claude.ai application and third-party platforms such as Amazon Bedrock or Google Cloud Vertex. It does not include Claude Code.
- ⁴ See our [Sonnet 3.7 report](#) and the Appendix.
- ⁵ We defined entrepreneurial activity as the share of conversations whose detailed request cluster is Entrepreneurship, Side income ideation, Capital raising, Creator monetization, E-commerce, Business models, Healthcare business, Social enterprise, or Event business. Resume activity are conversations where the artifact classifier classified the output as a resume or job application.
- ⁶ The time of day is based on inferring the state from the IP address of the conversation.

CHAPTER 2

Artifacts

In this chapter, we classify each conversation on chat and Cowork (hereafter “Claude conversations”)¹ by its artifact, which we sort into more than 30 categories. We refer to the primary output Claude produces in a conversation—a document, an explanation, a piece of code, an academic paper, and so on, whether presented in a chat window or as a separate document—as an artifact. The full list of artifacts is in the Appendix.

Our classifier identified 93% of Claude conversations as producing an artifact (Figure 2.1).² The most common artifacts are explanations (17% of conversations), documents and reports (15%), and guidance (11%). Conversational outputs (like explanations or guidance) and written deliverables (like documents or presentations) each account for about a third of conversations; code and technical work (like apps or scripts) for about a sixth.

What an output is doesn’t tell you what it’s for: the same artifact could be a work deliverable or a personal project. We look at that split next.




Figure 2.1: Claude’s outputs. The share of conversations with a specific output. The figure shows the twelve most common output types in Claude conversations. The 1P API mix is in the Appendix.

What is each artifact used for?

Our [January Economic Index](#) introduced a primitive that classifies each conversation as work, personal, or coursework. Here, we apply that split to the artifacts produced in Claude conversations (Figure 2.2).

Some categories of artifacts are almost always personal. More than 80% of conversations producing creative writing, guidance, and recipes were classified as personal. Within categories, the personal and work-related uses can look quite different. Personal creative writing, for instance, is dominated by fanfiction, worldbuilding, and poetry; the 13% that is work-related is mostly in the form of short-form video scripts, screenwriting, and speeches. Categories most likely to be work-related include creating marketing content (80%), creating blogs or articles (81%), and writing database queries (82%).

Many outputs are equally likely to be used for personal and work reasons, including creating plans or strategies (44% work-related, 49% personal) or translation (42% work, 44% personal). For example, the most common types of personal planning artifacts include travel itineraries and workout schedules, while work-related plans most often pertain to entrepreneurial or content strategies.

Finally, artifacts that are characteristic of coursework include creating academic papers and theses, educational materials and math-related queries, though a non-negligible share of each falls into both work and personal categories.

What each output is used for, Claude chat and Cowork

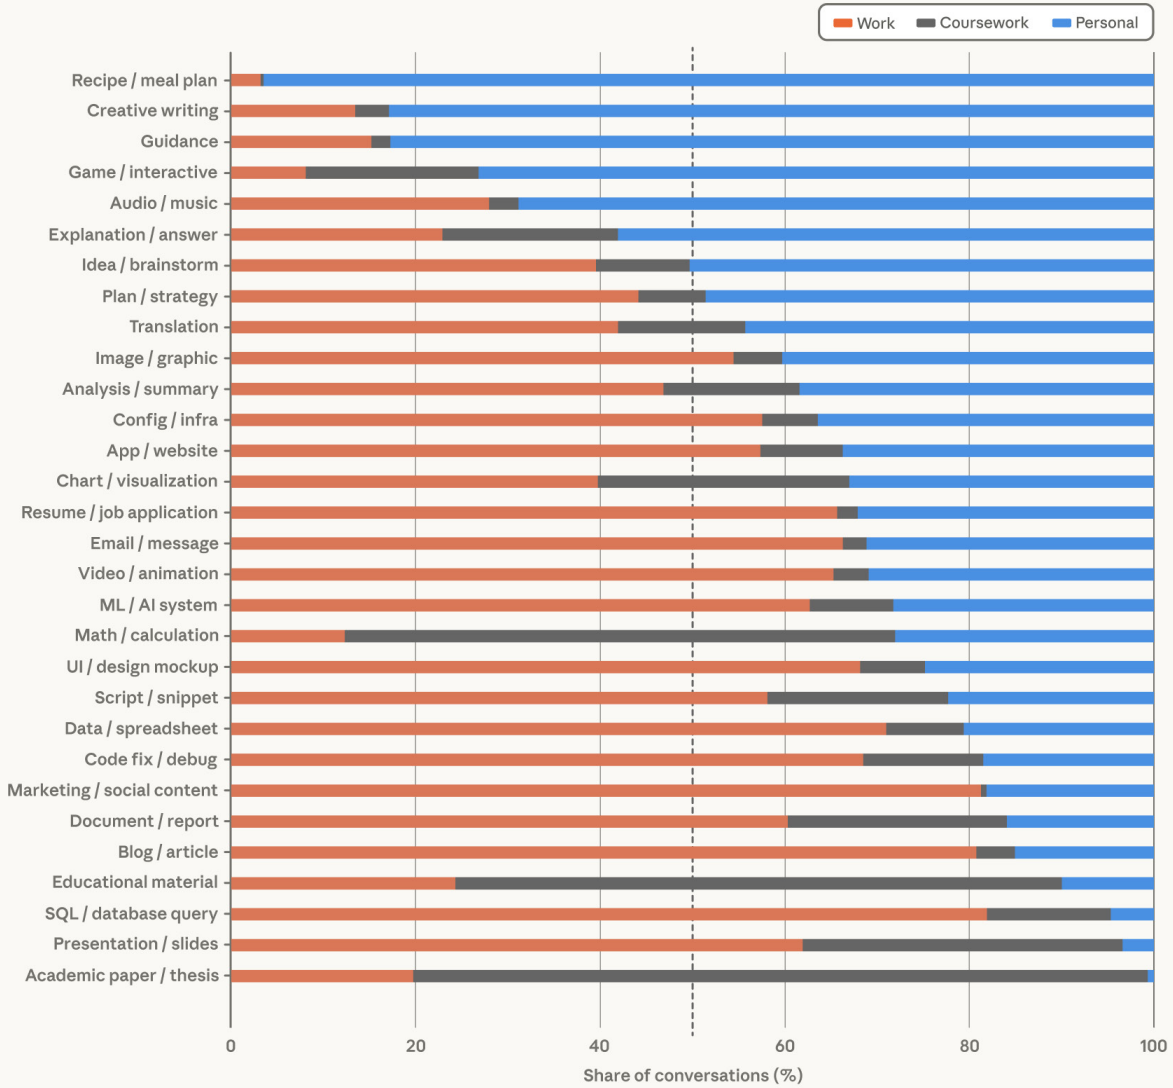


Figure 2.2: Share of artifacts by use case. Share of conversations with a specific output classified by its purpose into work, personal, and coursework.

We can also flip the question. Instead of asking what each output is used for, we can ask what sort of artifacts work, personal, and coursework conversations each tend to produce. Work conversations most often produce documents and reports (20%), followed by explanations (9%), email drafts (7%), and analyses and summaries (6%). Coursework conversations look broadly similar, with documents and reports leading there too (21%), closely followed by explanations (20%), educational materials (11%), and academic papers (6%). In contrast—and unsurprisingly—only 6% of personal conversations produce a document. Instead, the most common results are explanations (25%) and recommendations (22%).

Cost tracks the value of the work

Producing these outputs requires compute, and we find that compute tends to scale with the value of the work. We measure each conversation's computational costs in tokens—the amount of text processed and generated, including Claude's internal reasoning—and compare across occupations by mapping each conversation's classified task to the occupation that typically performs it. Throughout this section, we restrict our analysis to work-related conversations.

The left panel of Figure 2.3 shows a positive relationship between the median conversation-level number of tokens and the median wage in mapped occupation.³ For example, marketing managers earn roughly twice as much as editors (\$80 vs. \$37 per hour) and conversations mapping to their tasks consume approximately 2.5 times as many tokens. Admittedly, the relationship is noisy and there are notable outliers. Pharmacists, for example, earn nearly three times what statistical assistants do (\$68 vs. \$24 per hour), yet conversations mapped to pharmacist tasks use only about one twentieth as many tokens.

Token usage by occupation and by artifact

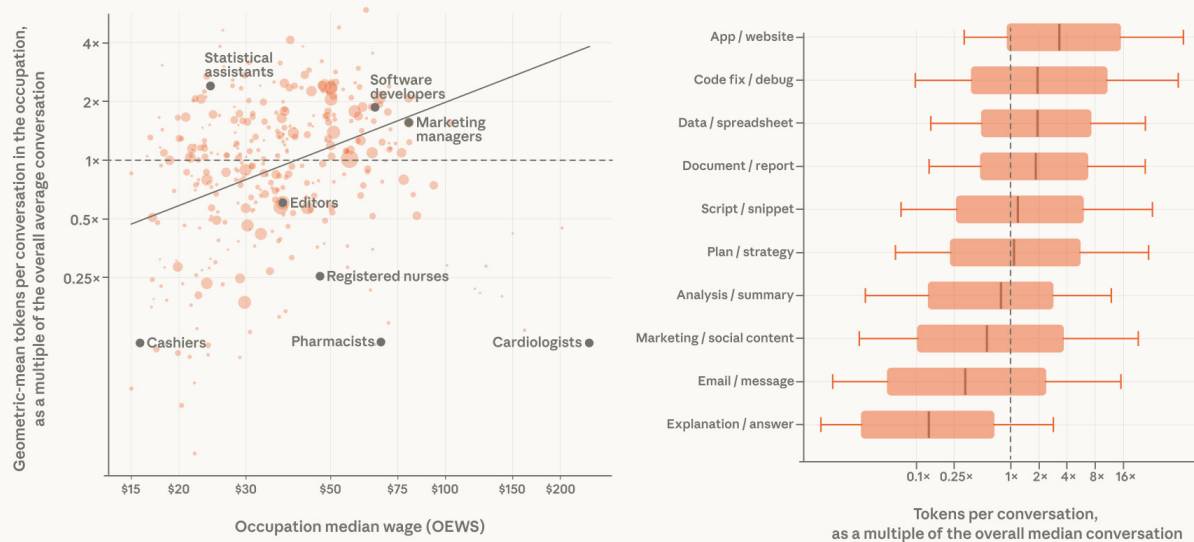


Figure 2.3: Conversations in higher-wage occupations consume more tokens. Left panel: Relationship between median wage for a given occupation and the typical (geometric mean) number of tokens in conversations classified to one of the tasks belonging to that occupation, normalized by overall mean. Both variables are represented on a logarithmic scale. Right panel: Distribution of tokens used to generate given artifacts. Black line illustrates the median, box represents p25 and p75, whiskers represent p10 and p90. All numbers are normalized by the overall median number of tokens and presented on the logarithmic scale. Data in both panels is restricted to chat and Cowork and restricted to conversations classified as work-related. Token counts are not adjusted for which model served the conversation.

The tokens consumed to generate different types of artifacts tell a similar story. More complicated and valuable outputs tend to consume significantly more tokens than simpler outputs. For example, conversations about building apps use more than three times the tokens of the median conversation. On the other end of the spectrum, a typical explanation uses about a fifth of the tokens of the median conversation. About 44% of the wage gradient in token consumption is explained by output mix—higher wage occupations are more likely to produce compute-intensive artifacts.

Why does this matter economically? In conversations mapped to higher-wage occupations, Claude produces more (1.34 times as much output per turn), while users engage more (1.53 times as many turns) and enable extended thinking more frequently (34% of conversations versus 31%; Table 2.4). Crucially, these move together: more production from Claude does not mean less from the user. If the human remains involved in the highest-value

tasks, the pattern looks more labor-augmenting than labor-displacing. It also shows that, to some extent, more valuable outputs cost more. The next section examines how much of the decision-making within each conversation is delegated to Claude.

Occupation median wage	Bottom third	Middle third	Top third
Tokens per conversation	1x	1.80x	2.07x
Turns per conversation	1x	1.58x	1.53x
Tokens per turn	1x	1.14x	1.35x
Claude's response per turn	1x	1.25x	1.34x
Price-weighted compute cost	1x	1.98x	2.05x
Extended thinking enabled (% of conversations)	31%	33%	34%

Table 2.4: What accounts for higher token consumption in higher-wage occupations. Features of conversations that contribute to higher token consumption in conversations mapped to higher-wage occupations. Occupations were grouped into terciles by their median wage, weighted by the number of conversations matched to each. Compute-related measures were normalized by their bottom-tercile's geometric means. E.g. the first row shows that a typical conversation mapped to a top-tercile occupation consumes 2.07 times as many tokens as a typical conversation mapped to a bottom-tercile occupation.

How much autonomy does Claude have to decide on its own?

We measure this on a 1-5 scale, from “none” to “extreme.” Tasks that are easy to describe or specify involve little autonomy: the lowest-autonomy outputs are math or calculations, translations, and Q&As. High-autonomy tasks are those that require selection among many possible choices, e.g., creating apps and websites, games, or presentations. Such work, which requires sustained judgment, has historically been difficult to automate. By comparing the level of autonomy in Claude chat and Cowork to Claude Code, we show that is starting to change.

Across almost all types of outputs (26 of 31 outputs shown) the level of AI autonomy is higher on Claude Code than chat or Cowork.⁴ For example, conversations producing scripts and code snippets involve 0.53 points more autonomy (on average, on the 1-5 scale) when created with Claude Code than conversations producing the same output on chat or Cowork. Across all conversations the average difference in autonomy is 0.37 points, and it has two main sources.⁵

Approximately two thirds of the difference is explained by the same tasks being executed with more delegation on Claude Code. Blog posts and articles illustrate this: the requests and tasks behind them are similar on the two surfaces, but the way people work with Claude differs sharply. The median chat and Cowork conversation producing a blog post or an article involves 13 rounds of back-and-forth, while the median blog-producing Claude Code session contains a single human prompt. The remaining third reflects the different mix of output types across the two surfaces.




Figure 2.5: AI Autonomy by output type. Average level of AI Autonomy by conversation output and surface. AI Autonomy is rated on a 1-5 scale from “none” to “extreme”.

One might suspect this difference simply reflects model choice. Claude Code sessions run on the most capable models far more often (54% are served by Opus, against 10% of chat and Cowork conversations). However, the gap persists when we compare conversations served by the same model. For example, among conversations using Sonnet, Claude Code sessions still show 0.26 points more autonomy, suggesting that the product used is likely more important than the underlying model.

Stepping back from the surface comparison, the output types where users delegate the most are the same ones that consume the most compute: across artifacts, mean autonomy and median token use rise together ($r = 0.68$ on chat and Cowork; Appendix Figure A.2).

Claude answers above the level it was asked

For each conversation, a classifier estimates two reading levels—one for the user’s prompt, one for Claude’s response—expressed as the years of education needed to understand the text.⁶ We find that reading level varies widely depending on artifact type. An average query resulting in an academic paper would require more than 16 years of education, roughly equivalent to bachelor’s level, and 15% of these conversations are at PhD level or above (20 or more years of education). On the other end of the spectrum are conversations resulting in recipes or guidance, where fewer than 10 years of education are required to understand the prompt.

In general, artifact types with higher-reading-level outputs also have higher-reading-level prompts (a correlation of 0.87 across conversations). However, we also observe that in almost every category, Claude’s output is at a higher comprehension level than the prompt, by roughly one year of education on average. The gap is widest where users describe something to be built, such as image and graphics (+2.6 years), games (+1.9), and apps and websites (+1.7). Some of the gap may simply be register; prompts are often terse and informal, while Claude tends to reply in polished prose. However, the gap is near zero for audience-facing writing (blogs -0.1 , academic papers $+0.0$, email $+0.3$), possibly because prompts typically draft language or source material written in the same register as the intended output.

Claude answers above the level it was asked




Figure 2.6: Reading level of user prompts and Claude's responses, by artifact. Estimated number of years of education needed to understand the user's prompt and Claude's response. Restricted to chat and Cowork.

-
- ¹ Data in this chapter cover chat and Cowork conversations sampled between April 10 and June 10, 2026. Where the autonomy discussion compares surfaces, Claude Code sessions from the same period are included. Wages are from the BLS OEWS, May 2025 release.
 - ² "None" is a catch-all for the conversations that didn't yield a prominent concrete output. This may include brief or abandoned exchanges, cases resulting in an error or cases where Claude asked a clarifying question and the user didn't continue.
 - ³ We use geometric means for the conversation level token counts since that variable is extremely right-skewed—a small number of conversations use several orders of magnitude more tokens than a “typical” conversation. The relationship is very similar if we use medians or if we weight the tokens by their respective cost to account for the mix of models used. There are some notable exceptions, including physician occupations.
 - ⁴ For a fuller picture of Claude Code, see our [companion report](#)
 - ⁵ The largest exception is data and spreadsheets, where Claude.ai conversations involve more autonomy than Claude Code (3.09 vs 2.74). This is mostly compositional: about 70% of the gap reflects a different mix of tasks. On Claude.ai this output leans toward financial modeling and dashboards, where Claude designs the structure; on Claude Code it leans toward structured extraction and tagging, where the specification is precise. Cowork, where data and spreadsheet work is both over-represented and especially autonomous, accounts for part of the Claude.ai lift.
 - ⁶ For details of the prompt used for this classification see Appendix to our [March report](#).

CHAPTER 3

Perceptions

The first two chapters show how people use Claude, but don't give much insight into the ways people experience AI at work—how they expect their jobs and workplaces to change, how they feel about AI's current and potential impact, and what they hope for from the technology. Our [interviews with 81,000 Claude users](#), conducted in December 2025 with [Anthropic Interviewer](#), [gave a picture](#): respondents reported large productivity gains, but also expressed worry about displacement. Those worries were concentrated among early-career workers and occupations where we observe Claude doing the most work.¹

In April 2026, we launched the [Anthropic Economic Index Survey](#) to build on this work. The survey allows us to ask people directly about their experience with AI and work, and to explore how responses vary with Claude usage. We link survey responses to usage data from mid-May to early June using [privacy-preserving methods](#). To characterize each respondent's usage patterns, we randomly sample up to 20 sessions per person within this time window (across Claude.ai, Cowork, and Claude Code, so that the mix of sessions reflects each person's typical usage across surfaces). We exclude respondents with fewer than five sessions to reduce sampling noise. Our final linked sample consists of about 9,700 survey respondents.

We find that most respondents expect significant AI progress over the next year. While people's perception of AI capabilities depends on their experience, where they live, and how exposed their job is to AI, their expectations about the pace of future progress are strikingly uniform, consistent with a [“rising tide,”](#) in which AI capabilities improve broadly.

Views on what that progress means for their own careers are less uniform. Early-career workers report that AI can do the highest share of their work and express the most concern about job loss. Yet—contrary to a common concern—the people who delegate to Claude the most are the most optimistic about their future labor market outcomes, and feel their skills are growing in value. And despite (or perhaps because of) their proximity to AI's frontier, the average respondent's hopes for the next decade center not on replacement but on collaboration. They hope AI can preserve meaningful work and automate the drudgery, and that its gains will be shared widely.

Who responded to the Economic Index Survey

The Economic Index Survey is not representative of the general population. We reach a random sample of Claude users, there may be selection in who completes the survey, and we filter out infrequent users from our analysis. Figure 3.1 shows the occupational mix of survey respondents (orange) alongside US employment (grey). Computer and Mathematical occupations are the most heavily over-represented, making up roughly 30% of survey respondents—comparable to their share of Claude usage, but far above their 4% share of US employment. Management, at 23% of respondents,² is also heavily over-represented relative to its 7% employment share, even though it accounts for only 4% of sessions. This gap is consistent with managers using Claude for tasks other than management itself: in the survey, judgment and management are named by many respondents (especially those with more experience) as capabilities AI lacks. Physical occupation categories like Transportation & Material Moving, Food Preparation & Serving Related, and Construction & Extraction are all under-represented in the survey, as they are in Claude sessions as well.

Survey occupational composition vs. US employment

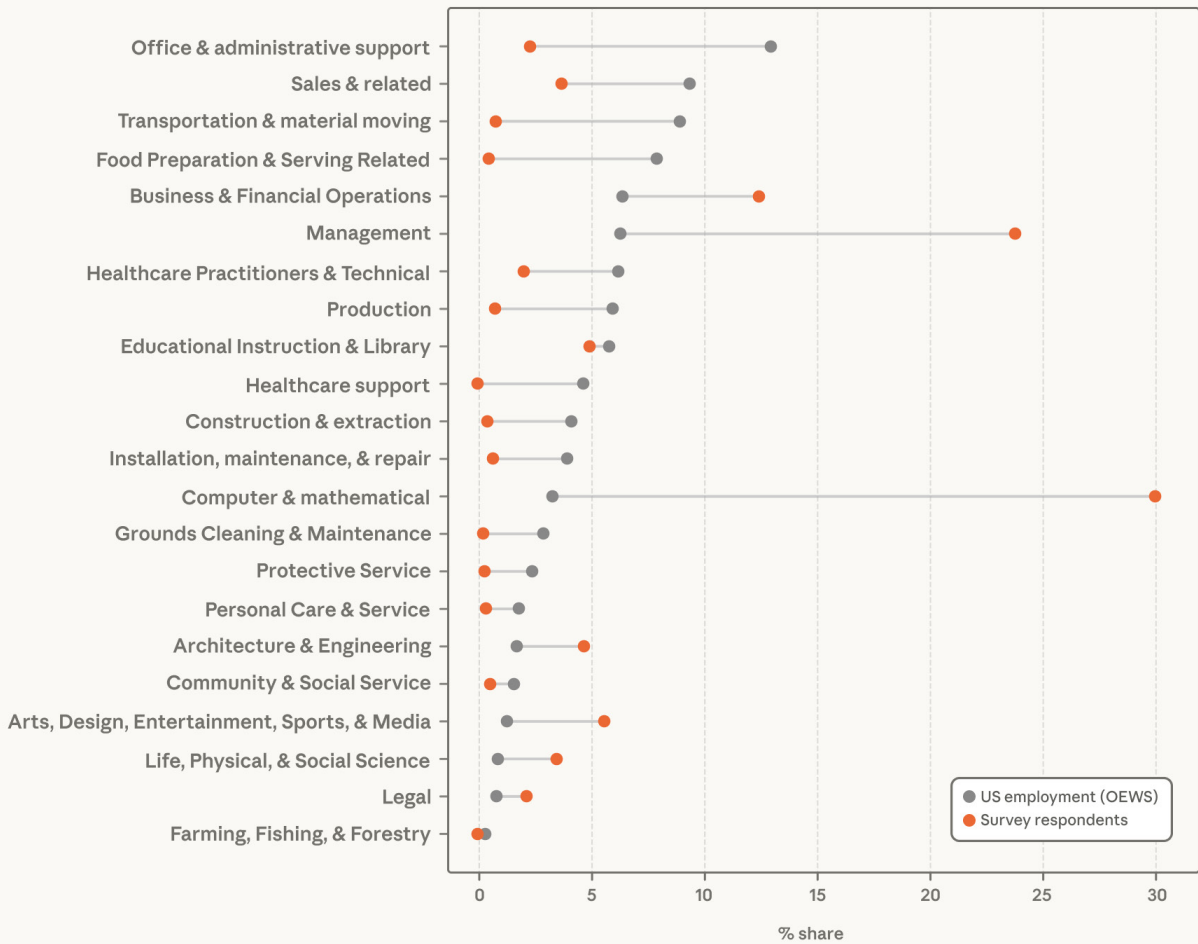


Figure 3.1: Survey respondents skew heavily toward computer & mathematical and management occupations relative to US employment. Share of survey respondents reporting their occupation in each of the 22 major SOC groups,³ compared to each occupation's share of US employment from OEWS.

AI and work tasks

Research on AI impacts often focuses on occupational exposure, or what share of tasks within a given job are doable with AI. In [prior work](#), we constructed a measure of *observed exposure*, which captures the share of occupational tasks we already see being done with Claude. We compared it to a commonly used measure of *theoretical exposure*, or the share of occupational tasks that a large language model could theoretically do.

Another way to understand occupational exposure is to simply ask people how much of their job AI is capable of doing. We asked respondents what share of their work tasks AI could do entirely on its own today (hereafter *reported exposure*), and what share they expect it to handle in 12 months (*anticipated exposure*), with the option to select from 5 bands ranging between “almost none” and “nearly all.” Close to 6 in 10 respondents chose a higher band for next year than for today. Over a third expect AI to be able to do most or nearly all of their work tasks next year (Figure 3.2).

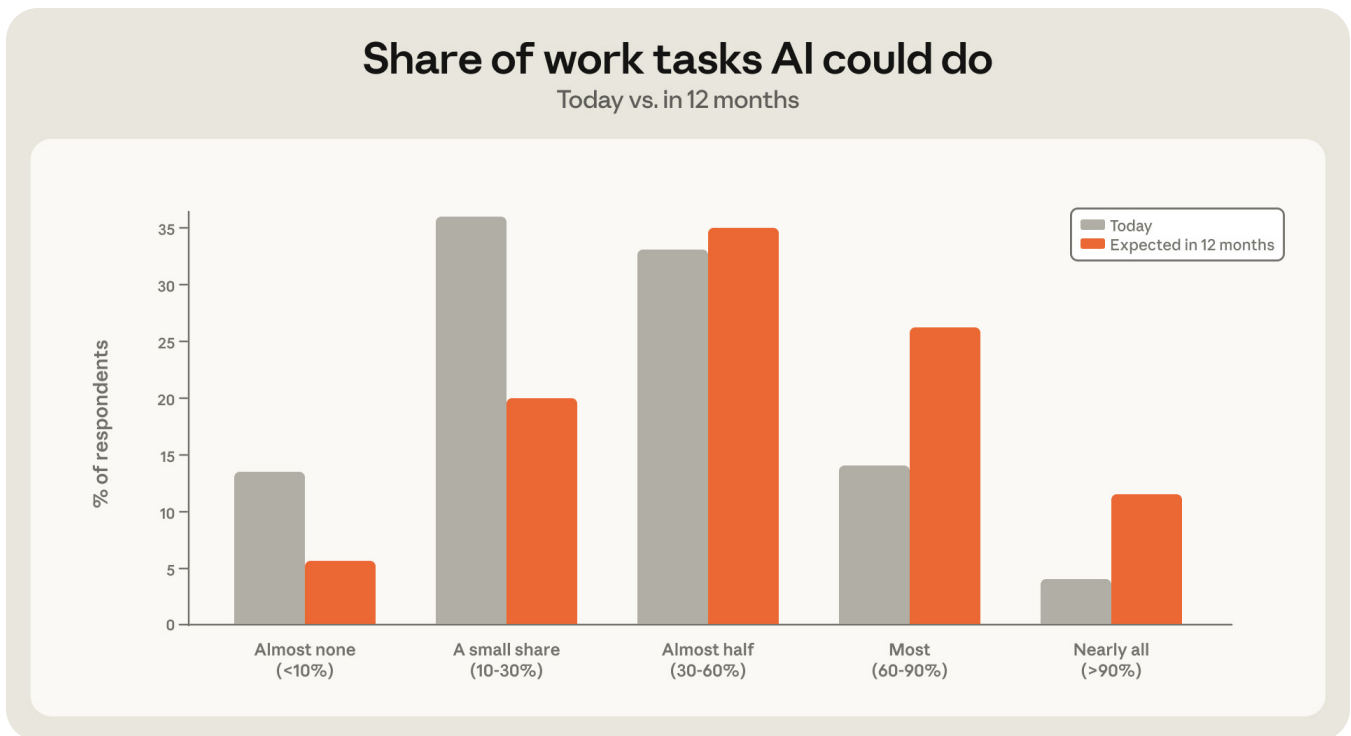


Figure 3.2: Most respondents expect the share of their work tasks AI could do to grow over the next 12 months. This figure shows the distribution of the share of their work tasks respondents say AI can do today versus in 12 months.

Figure 3.3 compares reported and anticipated exposure to observed and theoretical exposure. We ask whether what people report and anticipate AI can do lines up with the observed and theoretical exposure measures across occupations, and whether respondents whose occupations score higher on observed or theoretical exposure expect faster progress over the next year. On the first question, the answer is yes: reported exposure (grey dots) is positively correlated with both observed and theoretical exposure. On the second, the answer is no: the best-fit lines for reported and anticipated exposure 12 months from now (orange dots) are essentially parallel, meaning that people in roles

with high observed or theoretical exposure expect roughly the same increase in the share of their work tasks AI can do over the next year as those in roles with less observed and theoretical exposure.⁴ In other words, a software engineer and a construction manager anticipate roughly the same increment of progress within their profession.

It is also worth noting that reported exposure systematically exceeds observed exposure. One explanation for this is that not everybody does every task in an occupation, and our survey disproportionately reaches those who use AI more.⁵ Analogously, since theoretical exposure is an upper bound on what is possible instead of a measure of current use, theoretical exposure systematically overstates reported exposure.

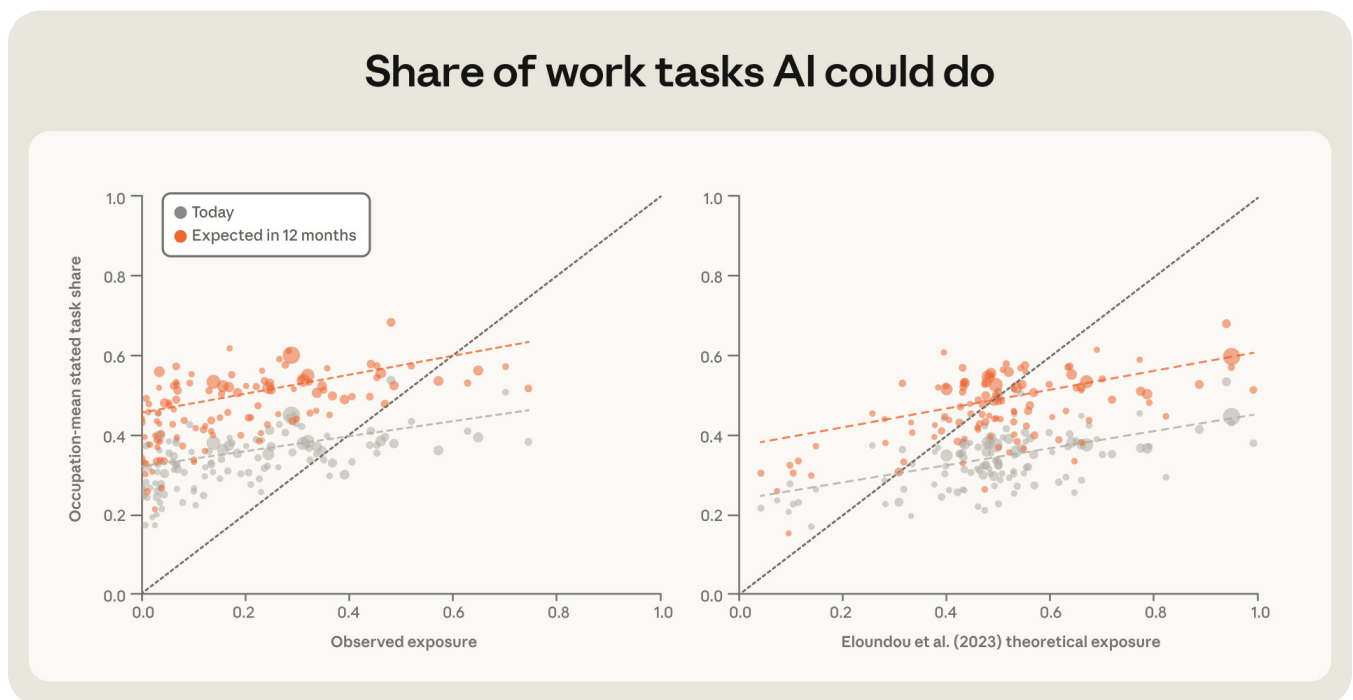


Figure 3.3: Reported and anticipated exposure vs other measures. The share of tasks people say AI could do today (reported exposure, in grey) and in 12 months (anticipated exposure, in orange) plotted against observed exposure (left panel) and theoretical exposure (right panel). Each point is an occupation, with occupations containing small samples grouped with exposure neighbors to ensure privacy. Reported and anticipated exposure are computed as the midpoint of the bin selected by the survey respondent.

We also examine how perceptions of AI’s current and future capabilities relate to the characteristics and usage patterns of respondents. The left panel of Figure 3.4 shows that perceptions of AI’s capabilities are negatively correlated with country GDP:⁶ The average share of tasks people report AI can do for them now is about 10 percentage points lower among high-income countries. This pattern is consistent with the possibility that AI substitutes for a larger share of the tasks that workers in lower-income countries do day-to-day, even if occupation-level exposure metrics—which tend to be higher in advanced economies—suggest otherwise. Indeed, [the IMF has noted](#) that while advanced economies face broader AI exposure overall, workers in lower-income countries may have less access to the complementary skills and infrastructure that allow AI to augment rather than replace their work. In [earlier work](#) we documented that lower-income economies tend to use Claude in more automated ways even when adjusting for differences in task mix.

The middle panel shows that reported and anticipated exposure are also negatively correlated with years of work experience.⁷ People with at least 15 years of experience put that share of tasks AI can do roughly 10 percentage points lower than those in their first year of work. We find evidence that this may be because experienced workers have accumulated tacit or context-specific expertise that is difficult for an AI to mimic. In follow-up questions, we asked people what tasks they thought AI would never be able to do and why; the most common responses emphasized that AI lacks the judgment, contextual awareness, and situational reasoning that their work requires. Respondents, and disproportionately those with at least 15 years of experience, also pointed to the relational and interpersonal dimensions of their jobs—building trust and managing people—as things AI cannot replicate.

As with occupational exposure to AI, we find that perceptions about future improvements in AI capabilities are essentially uncorrelated with GDP per capita and years of experience. The expected share of tasks that AI will be able to do in 12 months is uniformly higher than perceptions about AI’s capabilities today.

Reported and anticipated exposure

Relationship with people's characteristics and usage

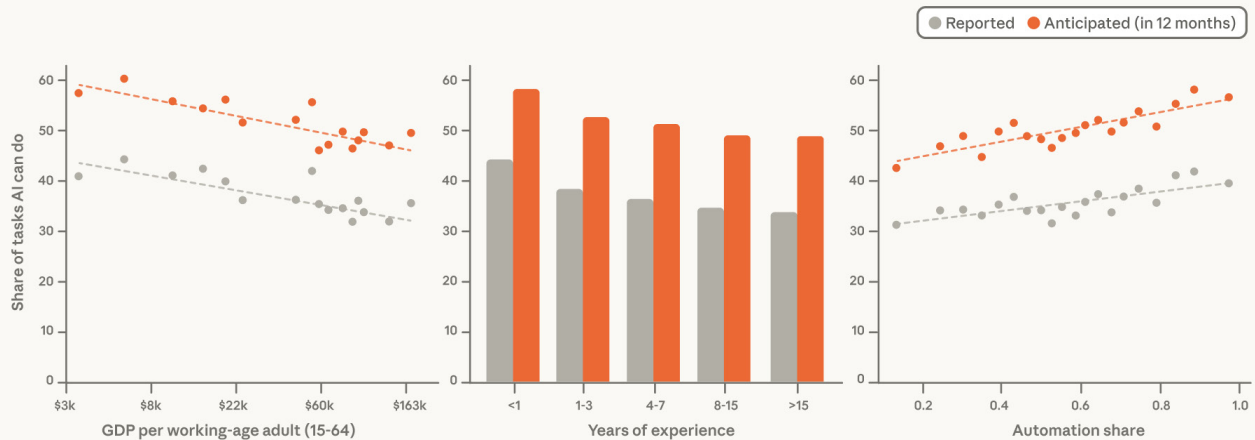


Figure 3.4: Reported and anticipated exposure are lower in higher-GDP countries and among more experienced workers, and rise with automated Claude use. Reported and anticipated (12 month) exposure against country GDP per working-age adult (left), years of experience (middle), and the share of Claude sessions classified as automated (right).

We next examine the relationship between how people interact with Claude and their current perceptions of Claude’s capabilities. As with past reports, we distinguish between “automation” and “augmentation” modes of collaborating with Claude. We identify conversations as automated when Claude is asked to complete a task with little to no input from the user. Concretely, automation share is the share of conversations whose pattern is either directive (“translate this document”) or a feedback loop (“edit this email...make it more casual”).⁸

The right panel of Figure 3.4 shows that reported and anticipated exposure rise with automation share. This could be because delegation is informative about capabilities—people who hand over entire tasks observe directly what AI can complete on its own—or because people who already believe AI can do their work are the most willing to hand it over. The same patterns hold when we replace automation share with the share of sessions devoted to work tasks, or the share conducted in Claude Code.⁹

AI and jobs

We also ask how people think their jobs will change in the next 12 months. More than a third of respondents said it was likely or very likely that responsibilities would significantly change (for themselves, a peer, a junior colleague, and a senior colleague). 10% rated losing their own jobs as likely or very likely. This is slightly below the annualized hazard rate of losing a job in the US;¹⁰ however, since our respondents skew toward knowledge workers in stable employment (a group that plausibly faces below-average separation risk at baseline), this may still indicate elevated perceived risk. When asked an open-ended question about what was driving their forecasts, 38% of the respondents who rated their job loss as likely or very likely attributed their forecasts to AI.¹¹ Notably, respondents were on average more worried about job loss for others than for themselves.¹² Respondents were especially worried about job loss for their junior colleagues, with over one third stating that the probability of a junior colleague losing their job in the next year was over 60%. Respondents were also more concerned about job loss (for everyone) in lower-income countries.

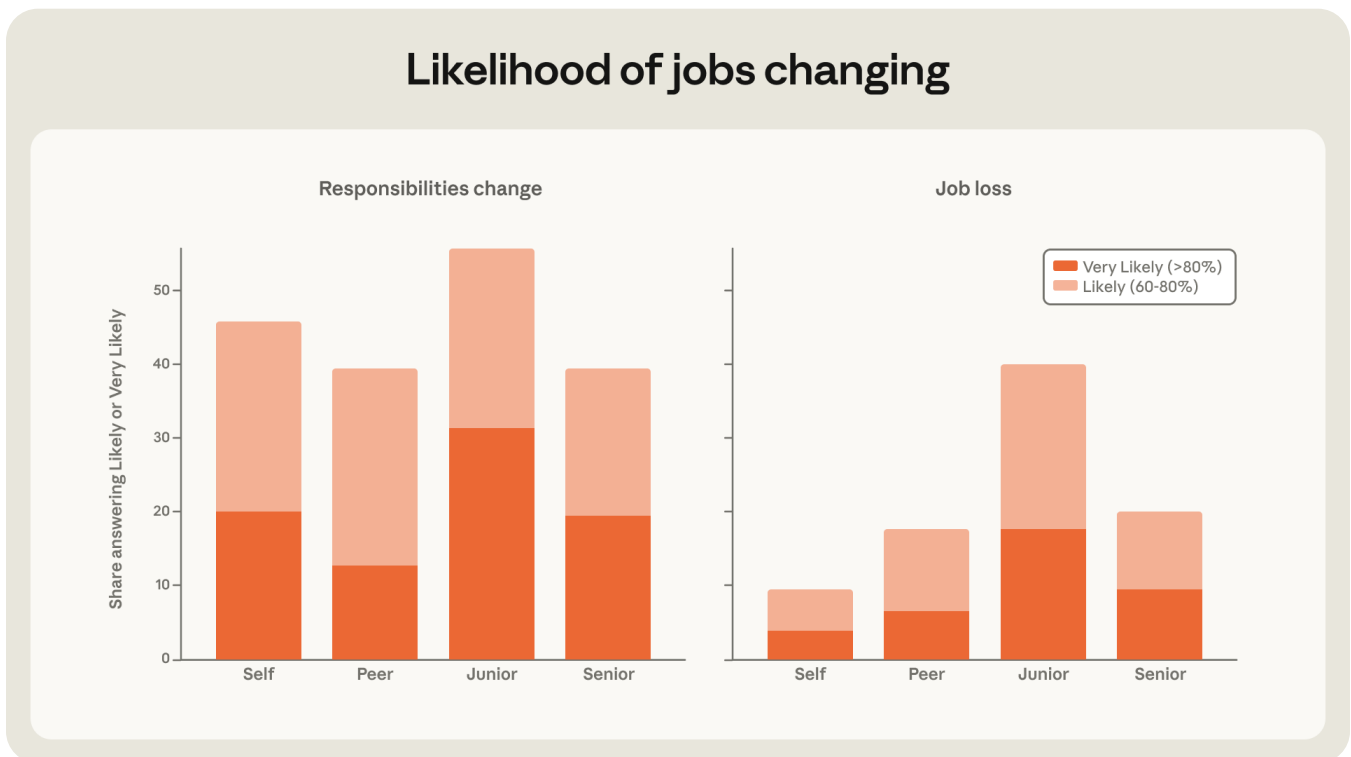


Figure 3.5: Share of people reporting it is likely or very likely job responsibilities will change significantly or people will involuntarily lose a job next year. Share of respondents that reported it was likely job responsibilities would change (left panel) or an individual would lose a job they wanted to keep (right panel) for themselves, a peer, a junior colleague, or a senior colleague.

Are people who use Claude in more automated ways also more worried about losing work? We examine what people said about AI’s expected impact over the next year on six dimensions of work: pay, job security, ability to find a new job (economic dimensions) and meaning, autonomy, and human interaction (intrinsic dimensions); and look at how these expectations differ by the automation share of Claude usage.

Across all six dimensions, people with a higher share of automated sessions feel more optimistic about the effect of AI on their job outcomes next year compared to those who use Claude more augmentatively. We saw the largest effects on expectations about positive impacts on future pay and ability to find a job.¹³

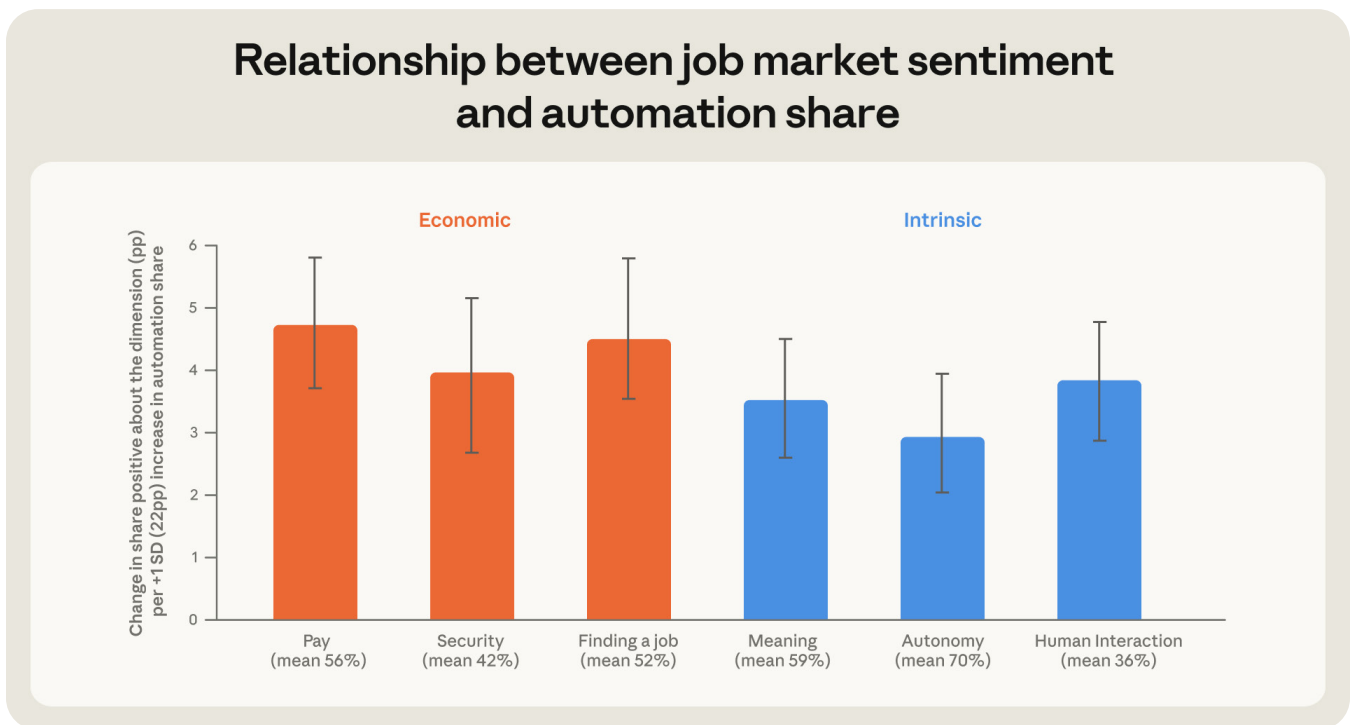


Figure 3.6: The share of people reporting positive expected impacts of AI on economic and intrinsic dimensions of job quality is higher among more automated Claude users. This figure shows coefficient estimates from linear regressions of an indicator for whether a person expected AI to have a positive effect on each of the six dimensions of job quality onto normalized AI automation share.

A natural question is why automated usage and sentiment move together. It’s possible that this relationship is explained by selection, that the people most enthusiastic about AI are also the most willing to hand over entire tasks to it. We can’t rule this out entirely, but these estimates don’t meaningfully change when we control for user tenure on [Claude.ai](https://claude.ai)—which we can think of as a proxy for enthusiasm, because it captures early versus later adopters.

Another possibility is that people who use AI in more automated ways experience more of its benefits today. Consistent with our [previous findings](#), large majorities of people report productivity gains in speed, scope, and quality of their work (86%, 82%, and 69%, respectively), while 27% report gains through cost savings on services they would otherwise have to purchase.

In addition to significant productivity gains, the majority of people also report learning more with AI (68%) and feeling like AI has made their skills more valuable (57%). Figure 3.7 shows how these two outcomes vary with the share of automated sessions. We see that the share of people reporting that AI is increasing the market value of their skills rises with automation share, while the share reporting they learn more is roughly flat.

A commonly voiced concern about delegation is that handing entire tasks to AI means offloading thinking, with gains in output coming at the cost of learning and skill atrophy. We do not see this pattern here: heavier delegators report learning at the same rate as everyone else. However, these are self-assessments, and skills can erode even as they become more valuable and as someone reports learning more, so the data do not rule out skill erosion.




Figure 3.7: The share of people reporting positive expected returns to their skills is increasing in automation share, while the share reporting they're learning more is flat. This figure shows the relationship between automation share of tasks and the share of people reporting AI increases the market value of their skills (blue) and they learn more when using AI (orange).

How usage differs between genders

So far we have explored how usage patterns relate to expectations and behavior. Next, we study who uses Claude in various ways. The most striking differences are by gender. Women, who make up only 12% of our linked respondent sample, use Claude differently from men. Even after accounting for occupational differences, they are marginally less likely to use Claude for work, their share of sessions in Claude Code is 0.24 standard deviations lower (6.3 percentage points), and their automation share is 0.33 standard deviations lower (7.3 percentage points). Instead, women tend to use Claude more iteratively, and they log more active time on chat than men, a signal of more collaborative engagement.¹⁴




Figure 3.8: Women have distinctly different usage patterns, even after conditioning on occupational differences. This figure shows women's usage patterns compared to men's. Each bar is the difference between women's and men's average on that usage measure, controlling for occupation (SOC minor groups), expressed in standard deviations of the outcome. Whiskers show 95% confidence intervals. The sample is restricted to respondents identifying as women or men. Women have significantly lower Claude Code and automation shares, while having higher total active minutes.

What do people hope for from an AI-transformed economy?

The Anthropic Economic Index Survey surfaces a mix of positive and negative experiences and sentiments with respect to AI, but we end the survey on a hopeful note. The final open-ended question asks respondents to “dream big: what do you hope an economy shaped by AI looks like in ten years?” We ran each survey response through a classifier which tagged responses with relevant themes. We show the top five most commonly cited themes below. Additional descriptions of each can be found in the Appendix.

The most common theme expressed was one of AI augmentation of work. Over half of survey respondents expressed some version of wanting to collaborate with AI on work that feels meaningful, of wanting their career to still matter, and/or hoping that new industries arise and create new job opportunities. Simultaneously, just over half of respondents hoped for AI automation—specifically of the tedious parts of their jobs—so they could have more free time and more space for meaning outside of work. The third most common theme, expressed by about one third of survey respondents, was one of shared prosperity: the hope that the economic gains from AI will be widely shared.

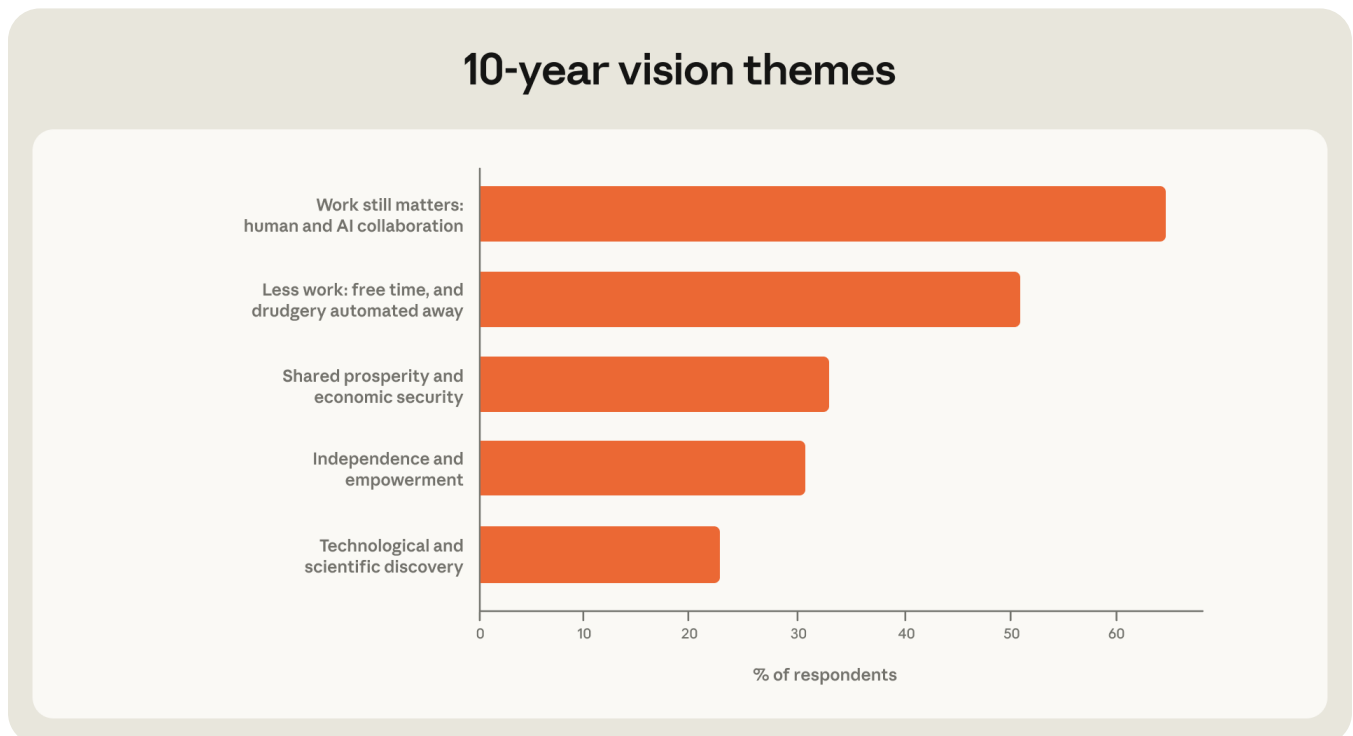


Figure 3.9: People hope for human-AI collaboration, automation of drudgery and more free time, and that economic gains from AI are broadly shared. This figure shows the five most common themes from people’s responses to an open-ended question on what they hope an AI-transformed economy looks like in 10 years.

-
- ¹ Similar patterns hold beyond our user base in more structured survey data produced by the [Anthropic Public Record](#), a nationally representative survey of more than 50,000 Americans.
 - ² Among people whose occupation was coded as management, 48.1% said they're employed at a company, 24.4% said they were a business owner with employees, and 21.7% said self-employed or contractor. The remainder are not currently employed, and are reporting their most recent occupation.
 - ³ Excludes military.
 - ⁴ Because responses are binned—so the lowest possible coded response exceeds zero, and the highest possible response falls short of one—the slopes in this figure are biased towards zero. As a result, we interpret the comparison of slopes qualitatively rather than as precise estimates. However, the patterns (positive slopes and close to parallel lines) are robust to instead estimating the relationship using an indicator for reporting that AI can do at least 60% of one's work tasks, which is unaffected by midpoint coding.
 - ⁵ The binned response scale likely also plays a role: because midpoint coding pulls reported task shares away from the extremes, observed exposure will tend to look as though it understates AI's capabilities in the least-exposed occupations and overstates them in the most-exposed ones, even absent any substantive difference.
 - ⁶ GDP per working-age adult is constructed in the same way as in the Anthropic Economic Index, using [World Bank WDI](#) (2024) and [UN World Population Prospects](#) (2024) for working age population estimates, and the [IMF World Economic Outlook](#) (2025 estimates) for GDP.
 - ⁷ We ask survey respondents about the number of years of experience they have in their current or closely related roles.
 - ⁸ The collaboration mode classifier maps transcripts to one of the following modes of interaction:
 - Directive: Human delegates complete task execution to AI with minimal interaction
 - Feedback Loop: Human and AI engage in iterative dialogue to complete task with human mainly providing feedback from the environment
 - Task Iteration: Human and AI engage in iterative dialogue to complete a task with the human refining the AI outputs
 - Learning: Human seeks understanding and explanation rather than direct task completion
 - Validation: Human uses AI to check or validate their own work
 - ⁹ Work share and Claude Code share are both positively correlated with automation: Claude Code is an agentic tool whose sessions are on average more automated than those on [Claude.ai](#), and work sessions likewise skew more automated than personal ones. Work usage also matters directly—the survey asks about work tasks, so people who use Claude for work may mechanically expect it to do a larger share of their work tasks. Conditioning on these measures therefore attenuates the relationship between automation share and task shares (today, in 12 months, and the change), but all three relationships remain positive and statistically significant.
 - ¹⁰ The US layoffs and discharges rate ([BLS JOLTS](#), total nonfarm, seasonally adjusted) averaged ~1.1% of employment per month over the 12 months through April 2026, amounting to a ~13.4% annualized sum, so 10% is slightly below the realized annual incidence of involuntary separation events.
 - ¹¹ This question was asked about the job change forecast and job loss forecast together. The 38% is therefore an upper bound on the share of people who attribute their own job loss forecast to AI.
 - ¹² This mirrors a familiar pattern of people rating their own circumstances more favorably than other people's. A similar phenomenon was [observed during Covid](#), when self-reported financial well-being exceeded perceptions about the national economy.
 - ¹³ This contrasts with country GDP and experience, where lower GDP and experience correlate with higher task shares (as shown in Figure 3.4) and higher stated job loss probabilities.
 - ¹⁴ While this could be due to substitution between chat/Cowork and Claude Code, the pattern holds even when controlling for Claude Code session share. These patterns also survive controlling for occupation fixed effects.

Discussion

AI is diffusing rapidly throughout the economy, across an increasing number of surfaces, with increasingly intelligent outputs. In earlier AI chat interfaces, usage was simple, contained in the chat window without web search, tool calls, artifacts, or other affordances. Now, Claude models can operate autonomously for hours through Claude Code and Cowork. As these forms change, the user base is shifting as well. Early adopters were highly technical. Our most recent users apply Claude to tasks that command [lower wages](#) in the labor market.

In this report, we took several steps toward more informative measurement. First, we began measuring more and more frequently, processing data in hourly samples. This reveals how the cadences of daily life are etched into our usage logs and opens avenues for future research. Second, we began recording artifacts, or the outputs that people take away from Claude. These make Claude's work output more legible, and show some intuitive patterns.

Finally, usage data only carries so much information. Our survey allowed us, for the first time, to ask people directly about how they use AI and what they feel about it. We found that our survey respondents use AI for more than we give it credit for—they report AI can do a higher share of their work than the observed exposure measure for their occupation would suggest. Asked to forecast next year's capabilities, over 35% predicted that AI would be able to do most of their work.

Accurately classifying the work that Claude does will remain a moving target. For example, as AI capabilities increase, AIs may increasingly interact and exchange with each other, perhaps in ways inscrutable to humans or simple classifiers. Ultimately, Claude's impact on the economy will be visible in economic aggregates like [employment](#) and productivity as much as its usage logs. Still, AI is likely to have its earliest impacts in the areas where it does the most work, so shedding light on these ever-changing usage patterns will remain a key way to inform the public.

Appendix

Available [here](#)

Citation

```
@online{anthropic2026aeiv6,  
  author = {Maxim Massenkoff and Eva Lyubich and Szymon Sacher and  
Zoe Hitzig and Shaoyi Zhang and Ryan Heller and Peter McCrory},  
  title = {Anthropic Economic Index report: Cadences},  
  date = {2026-06-26},  
  year = {2026},  
  url = {https://www.anthropic.com/research/economic-index-june-2026-  
report},  
}
```