

ANTHROPIC

Appendix to “The AI Expectations Gap”

October 2026

A. Survey samples

The survey of the general public was fielded by Morning Consult between August 11 and 23, 2026, and collected 16,395 completed interviews. Some respondents appeared multiple times and we included only their first interview which left us with 10,980 respondents. Each of the eight daily samples was weighted to the US adult population and, after removing repeat observations, we recomputed those weights so that our sample reproduces the full sample weighted by age and gender, race and ethnicity, education, census region, and household income. Several questions were asked randomly about 2027 and 2030, each to a half of the sample. The staff poll ran from August 9 to 12, 2026, and drew 1,334 respondents and every department of the company was represented.

In the general population survey respondents could select “not sure” (chosen by between 12 and 21%); we exclude these question-by-question. The Anthropic staff poll did not include the “not sure” option, instead respondents could skip questions. When we map survey responses to model-implied forecasts, we use only the respondents who completed every required question: 3,259 US adults (only the 2030 half) and 1,093 staff.

B. Which beliefs about AI are responsible for differences in model-implied growth

How much of the gap in model-implied GDP growth between US adults (5.4%) and Anthropic staff (17.1%) does each belief account for? To find out, we start from the answers of US adults and swap in the staff’s answers one at a time. We swap US adults’ responses with staff’s responses at equivalent percentile of the response (breaking ties randomly). The impact of swapping a specific answer depends on respondents’ other answers. For example, the impact of adoption depends on AI’s capabilities. Suppose only adoption and capability answers were swapped. If capability is swapped first, it adds 1.2 percentage points of 2030 growth. If adoption is swapped first and capability second, then capability adds 2.0 points. We therefore run the model for all combinations of swaps and

average each belief's contribution over the 120 possible orders. Figure A1 shows how much each contribution varies depending on the order.




Figure A1: How much each belief adds, across all orders. Contribution of each of the staff's beliefs to the change in the median implied 2030 GDP growth of US adults when all five of the staff's answers are swapped in. Bars show the average contribution over all 120 orders of swapping. Lines show the smallest and largest contribution across those orders. They are not confidence intervals.

C. Additional Figures

The figure below shows the full distributions of the answers behind Figure 4.




Figure A2: Time a task takes with AI compared with without it. Respondents were asked how long a task that today's AI is well suited for takes with AI, counting the time spent checking and fixing its work, compared with not using AI. Segments show the share giving each of the six answers, from "about a tenth of the time or less" (left) to "it takes longer than without AI" (right).

D. External validation against the Anthropic Economic Index

The survey's measure of AI use lines up with observed usage. Across states, the share of adults who used AI for work in the past three months rises with the state's usage index in the Anthropic Economic Index. The index is the state's share of US Claude.ai usage divided by its share of the working-age population. Weighting states by their number of respondents, one unit of the index goes with 8.1 more percentage points of adults using AI at work (standard error 2.0, correlation 0.57). The one outlier is the District of Columbia. Its high index value may reflect commuters' daytime usage (Figure A3).

AI use at work by state, survey and Economic Index




Figure A3: Share of adults who used AI for work, by US state, against the Anthropic Economic Index. The vertical axis is the share of the state's survey respondents who used an AI tool for a work task in the past three months. The horizontal axis is the state's usage index in the [Anthropic Economic Index](#): its share of US Claude.ai usage divided by its share of the US working-age population. Marker area is proportional to the number of respondents in the state. The District of Columbia (index 3.3) lies beyond the right edge and is included in the fit.

The quiz website records the country from which the visitor is connecting. We check how this correlates with country-level AI usage in the Anthropic Economic Index. Visitors came from 146 countries. Across the 45 countries with at least 50 respondents, the number of responses per working-age person correlates strongly with the country's usage index ($r=0.87$ in logs). The median growth implied by a country's respondents is unrelated to its usage index, and differences between countries account for less than 1% of the variation across visitors. The median US visitor's answers imply 12.2% growth, against 12.7% for visitors elsewhere (Figure A4).

Website visitors by country




Figure A4: Quiz responses per working-age person against the Anthropic Economic Index, by country. Each point is one of the 45 countries with at least 50 website visitors who completed the Explorer's quiz in the 48 hours after its launch. The vertical axis is a country's share of those visitors divided by its share of the world's working-age population, so that 1 is the world average. The horizontal axis is the country's usage index in the Anthropic Economic Index. Both axes are on log scales. Marker area is proportional to the number of respondents.