

Technical Appendix

Data sources

This analysis is based on three data sources:

- Work requirements and job skills data from the U.S. Department of Labor's Occupational Information Network (O*NET, Version 30.3)
- Data on the occupation and employment of workers from the Current Population Survey (CPS)
- Occupational employment estimates from the Bureau of Labor Statistics' Occupational Employment and Wage Statistics ([OEWS](#), May 2023)

We used CPS data from 2017 through 2019 and January 2023 through June 2026. We omitted 2020 through 2022 because the COVID-19 pandemic disrupted employment patterns. This analysis includes about 4 million records covering 480 occupation groups.

AI exposure measure

We used a method developed in a [2023 Pew Research Center study](#) and updated it in 2026 with newer O*NET 30.3 data on work activities and new activity classifications.

- We identified some work activities that artificial intelligence can perform or assist with and others that it can only do minimally. We then scored each detailed occupation based on how important those activities are to it. (Read more on this below.)
- We combined detailed occupations into CPS occupation groups, weighting each occupation by its employment. Detailed occupations were matched to CPS occupation codes using the BLS National Employment Matrix crosswalk.
- We normalized the index so the least AI-exposed occupation in our data gets a score of zero. That does not mean an occupation has no exposure to AI — only that it is the lowest-scoring occupation among those we measured.
- The scores are continuous. We reported the figures as the effect of a 1-point increase, computed as $(e^{\beta} - 1) \times 100$.

How we determined which work activities are most and least exposed to AI

We classified O*NET's 41 work activities by how exposed they are to generative AI:

- 16 activities centered on reading, analyzing, writing, communicating and advising were classified as having high exposure.
- Nine activities requiring a physical or in-person presence were classified as having low exposure.
- The remaining 16 activities combine both kinds of work and were not included in the score.

Each occupation's AI exposure score is the average importance of its high-exposure activities minus the average importance of its low-exposure activities. Occupations that did not include any high- or low-exposure activities (i.e., had no score) were assigned the median score.

Work activities with high exposure to AI

Activity ID	Activity
4.A.1.a.1	Getting information
4.A.2.a.4	Analyzing data or information
4.A.4.b.3	Training and teaching others
4.A.4.a.3	Communicating with people outside the organization
4.A.4.b.6	Providing consultation and advice to others
4.A.2.a.3	Evaluating information to determine compliance with standards
4.A.2.b.2	Thinking creatively
4.A.3.b.1	Working with computers
4.A.2.b.1	Making decisions and solving problems
4.A.3.b.6	Documenting/recording information
4.A.2.b.4	Developing objectives and strategies
4.A.4.a.2	Communicating with supervisors, peers or subordinates
4.A.4.a.1	Interpreting the meaning of Information for others
4.A.2.b.3	Updating and using relevant knowledge
4.A.2.b.6	Organizing, planning and prioritizing work
4.A.2.a.2	Processing information

Note: The activity ID is the identifier for a work activity as listed in the O*NET data.
Source: Pew Research Center analysis of O*NET (Version 30.3).

PEW RESEARCH CENTER

Work activities with low exposure to AI

Activity ID	Activity
4.A.3.b.4	Repairing and maintaining mechanical equipment
4.A.1.b.2	Inspecting equipment, structures or materials
4.A.3.a.2	Handling and moving objects
4.A.4.a.5	Assisting and caring for others
4.A.3.a.4	Operating vehicles, mechanized devices or equipment
4.A.3.a.1	Performing general physical activities
4.A.3.b.5	Repairing and maintaining electronic equipment
4.A.4.a.8	Performing for or working directly with the public
4.A.3.a.3	Controlling machines and processes

Note: The activity ID is the identifier for a work activity as listed in the O*NET data.
Source: Pew Research Center analysis of O*NET (Version 30.3).

PEW RESEARCH CENTER

Our statistical model

The model tests whether women and men in the same occupation experienced similar employment growth. It also looks at whether any differences become larger in occupations with greater AI exposure.

$$\log(\text{Employment}) = \beta_0(\text{AI Exposure} \times \text{Post-2023}) + \beta_1(\text{Female} \times \text{AI Exposure}) + \beta_2(\text{Female} \times \text{Post-2023}) + \beta_3(\text{Female} \times \text{AI Exposure} \times \text{Post-2023}) + \beta_4(\text{Female}) + \alpha(\text{Occupation}) + \gamma(\text{Year}) + \varepsilon$$

The model separates out anything common to a specific occupation or a specific year, so what's left reflects only the part of the AI-exposure story that's specific to men or women after 2023.

- β_0 is the coefficient that measures how much AI exposure predicts employment growth after 2023 for men.
- β_3 is the coefficient that measures how much steeper (or flatter) the relationship between AI exposure and employment growth is after 2023 than before 2023 for women compared with men. Women's total post-2023 coefficient (on the log scale) is $\beta_0 + \beta_3$.
- The regression uses employment counts by occupation, year and sex, not individual records. Standard errors are clustered by occupation.

We weighted the regression using each occupation's average employment level from before 2020. This means occupations that employed more people before the pandemic had proportionally more influence on the estimated gender gap than occupations with fewer people did.

We kept the weighting constant across the whole analysis period, so an occupation's weight wouldn't shift over time as its employment grew or shrank. We did this because the analysis is testing whether AI exposure influenced growth. If we let growth determine the relative importance of each occupation, we'd be using the outcome we're measuring to help decide how much that outcome counts. Using pre-pandemic employment avoids that: Each occupation's weight is locked in before the growth we're studying even happens, so it can't feed back into the result.

Alternative measures of AI exposure

We tested our findings against four other published measures of AI exposure. One was based on how people [actually use Anthropic’s AI model, Claude, at work](#), not just what tasks AI can do. The others link AI capabilities to the abilities each occupation requires — one broadly across [10 AI capabilities](#), and two narrower versions that isolate exposure to [language modeling](#) and [image generation](#). All five estimates use the same version of the model, which also accounts for differences specific to each occupation in each year.

How the gender gap holds up under different ways of measuring AI exposure

% gender gap in employment growth, shown per point and per standard deviation of exposure for each measure

AI exposure index	Per 1-point increase in AI exposure*	Per S.D.	Type of exposure **
Pew-derived (updated to O*NET 30.3)	-6%	-5%	Theoretical
Anthropic observed exposure	-28%	-5%	Observed
AI Occupational Exposure (Felten-Raj-Seamans)	-7%	-7%	Theoretical
Language Modeling AIOE	-8%	-7%	Theoretical
Image Generation AIOE	No notable differences	No notable differences	Theoretical

* The percentages in the “Per 1-point increase in AI exposure” column are not comparable across rows because each index uses a different scale. The percentages in the “Per S.D.” column are comparable.

** “Theoretical” exposure estimates which tasks AI could perform based on task content. “Observed” exposure is based on how people use Anthropic’s Claude at work.

PEW RESEARCH CENTER

Each index uses its own scale for AI exposure, so a 1-point increase means something different on each index. That means the “Per 1-point increase in AI exposure” column isn’t directly comparable across rows. The “Per S.D.” column puts every index on the same footing by measuring the effect of a 1-standard-deviation increase. This column is comparable across all five measures.

Across four of the five measures, the gender gap stayed between 5% and 7% per standard deviation of exposure.

The Image Generation AIOE was the one exception: Its estimated gap was smaller and not statistically distinguishable from zero. This is expected because image generation tools have little to do with the reading, writing and information tasks that make up most AI-exposed work, so this measure serves as a check.