Washington Center for Equitable Growth
Evidence for a stronger economy

equitablegrowth.org

Working paper series

AI exposure by U.S. occupations and work tasks and the effect on wages

Chiara Chanoi
Chris Bangert-Drowns

October 2025

https://equitablegrowth.org/working-papers/ai-exposure-by-u-s-occupations-and-work-tasks-and-the-effect-on-wages/

facebook.com/equitablegrowth
@equitablegrowth
equitablegrowth.org/feed
info@equitablegrowth.org
740 15th St. NW 8th Floor
Washington DC 20005
202-545-6002

AI exposure by U.S. occupations and work tasks and the effect on wages

By Chiara Chanoi and Chris Bangert-Drowns

October 2025

Executive summary

This analysis of labor market AI exposure builds on previous work from the Pew Research Center and the AI firm Anthropic as well as Equitable Growth's own job quality series, confirming differences in AI exposure by gender, race, education, and income. Women tend to work in more exposed occupations compared to men, while White and especially Asian workers tend to work in more exposed occupations compared to other racial groups. Exposure is larger for people who work high-paying, high-education jobs, regardless of gender or race. Our research finds a small but statistically meaningful positive correlation between AI exposure and income. It additionally appears that augmentative uses of AI are associated with higher wages while automative uses are associated with lower wages.

Terms & definitions

Exposure: The share of total AI-induced work task changes that fall on a particular occupation or group of people. AI exposure in this analysis refers to the share of Claude.ai queries associated with a particular work task. This task-level AI exposure is aggregated at the work activity and occupation levels and combined with public survey data to estimate AI exposure by race, gender, and other demographic characteristics.

Augmentation: AI-induced work task changes that do not replace human labor but do make work more efficient. Anthropic defines augmentation as when "AI *collaborates* with a user to perform a task" (emphasis added) and lists three types of augmentative uses: validation (defined as work verification and improvement), task iteration (collaborative refinement process), and learning (knowledge acquisition and understanding.

Automation: AI-induced work task changes that directly replace human labor. Anthropic lists two types of automative AI uses: feedback loop (task completion guided by environmental feedback) and directive (complete task delegation with minimal interaction).

Occupation: Particular job types as defined by the O*NET and Current Population Survey datasets. The O*NET set of occupations is more detailed than the CPS set, with some CPS occupations broken into several O*NET entries.

Worker tasks: Functions performed by a worker in the course of doing a particular job. The O*NET database links each work task with only one occupation. Some tasks are identical or very similar across occupations. For example, a task assigned to postsecondary Business teachers is to “Evaluate and grade students’ class work, assignments, and papers.” A postsecondary Computer Science teacher is assigned a comparable task: “Evaluate and grade students' class work, laboratory work, assignments, and papers.” However, most tasks are unique to their linked profession. For example, a task for bartenders is to “create drink recipes”.

Work activities: An intermediate stage in the O*NET hierarchy, activities refer to broader functions that combine various work tasks and can be performed across various occupations. For example, the bartender-specific task of “creating drink recipes” is sorted into the broader work activity, “Thinking Creatively.” Conversely, the barber’s task, “stay informed of the latest styles and hair care techniques” qualifies as part of the work activity “Updating and Using Relevant Knowledge”. O*NET also produces measures of the relative importance of a work activity to each occupation.

Executive summary

This analysis of labor market AI exposure builds on previous work from the Pew Research Center and the AI firm Anthropic as well as Equitable Growth’s own job quality series, confirming differences in AI exposure by gender, race, education, and income. Women tend to work in more exposed occupations compared to men, while White and especially Asian workers tend to work in more exposed occupations compared to other racial groups. Exposure is larger for people who work high-paying, high-education jobs, regardless of gender or race. Our research finds a small but statistically meaningful positive correlation between AI exposure and income. It additionally appears that augmentative uses of AI are associated with higher wages while automative uses are associated with lower wages.

SECTION 1: Introduction

Recent advancements to Artificial Intelligence (AI) – particularly to chatbots such as Claude.ai and ChatGPT – have made waves in the labor market. The cognitive power of artificial intelligence has the potential to enhance the productivity of some workers while automating away certain worker tasks. The progression of AI in the workplace is therefore of scientific and policy interest. By utilizing federal employment data and publicly available AI use measurements, Equitable Growth is able to measure worker sensitivity to artificial intelligence and assess how such exposure may affect wages. We find that individuals in high-paying, high education jobs are more likely to encounter AI in the workplace, and verify AI as a significant but minor positive predictor of wages. However, such exposure rates are modified by other worker demographics such as race, gender, and age.

Additionally, AI's effect on income, while overall marginally positive, may be contingent on how it is used in the workplace.

The extent to which workers are affected by artificial intelligence is known as worker exposure. It is typically theorized to manifest in one or both of two forms:

- Augmentative exposure, in which AI supplements worker responsibilities, and
- Automative exposure, in which AI performs the work.

The Pew Research Center and AI firm Anthropic have both assessed the growth of AI use in the workplace. By determining which worker tasks and activities are commonly supplemented or replaced by AI, they calculate AI exposure by occupation type and worker demographics. Both utilize the work activity hierarchy supplied by the U.S. Department of Labor's Occupational Information Network (O*NET) as the framework for their studies. O*NET offers particular insight into how workers are currently exposed to professional changes brought on by AI. The detailed data have historically been useful for understanding job content and is the basis for Equitable Growth's latest series on job characteristics and job quality in the United States.

O*NET identifies 41 work activities, or actions that generally occur as part of work. These generalized activities are broken down into more specialized "intermediate," then "detailed" work activities, finally fragmenting into approximately 20,000 worker tasks that align specifically with certain occupations. In 2023, the Pew Center engaged a team of experts to estimate which of O*NET's work activities might be highly, moderately, or minimally exposed to AI. They then used these work activity exposure levels to determine which professions – and professionals – were susceptible to increased AI use. In 2025, Anthropic connected queries made to their chatbot, Claude.ai, to certain O*NET worker tasks and their corresponding occupations. These measurements are some of the first to offer calculable data on AI and how it may affect the workplace.

By aggregating Anthropic's task exposure metrics to the occupational level, then matching these vocations to Census employment data, we quantifiably update the Pew Center's 2023 findings. We find that while work activities are exposed to AI at different rates than anticipated by Pew, much of our occupational and demographic findings are consistent with its expectations. Women tend to work in more exposed occupations compared to men, while White and especially Asian workers tend to work in more exposed occupations compared to other racial groups. Exposure is larger for people who work jobs requiring greater education and higher paying wages, regardless of gender or race. In addition to calculating these expanded statistics on AI exposure, we are also able to combine AI use rates by occupation with Census income data. This provides the basis for analysis of the current relationship between AI exposure and wages. **We find that AI exposure has a statistically significant and positive but minor effect on hourly wages, which may be due to more profound but opposing influences of different types of AI.** Our results indicate that higher automative exposure is consistent with lower wages, while higher augmentative AI exposure predicts higher hourly rates.

SECTION 2: Pew Occupational exposure

In July 2023, Pew published its paper on AI exposure finding higher exposure among women, Asian, college-educated, and higher-paid workers. Pew's methodology was largely qualitative, with the researchers using their "collective judgment" to assign different levels of exposure to the various work activities in the O*NET database. The Pew researchers then aggregated work activities to the occupation level, using importance ratings data from O*NET to produce an occupation-level measure of relative AI exposure. Occupations were then ranked by exposure, with the top 25% labeled high exposure and the bottom 25% labeled low exposure.

The Pew researchers linked their O*NET-derived dataset to the BLS's Current Population Survey and produced a set of summary statistics estimating variation in AI exposure across a range of demographic characteristics. Overall, 19% of workers were employed in jobs with high AI exposure, compared to 23% in low-exposure jobs. Women were more exposed than men, though even among women a slightly larger share of workers were employed in low-exposure jobs. White and Asian workers were more likely to work in high-exposure jobs, while Black and especially Hispanic workers tended to work in low-exposure jobs.

Possibly the most important driver of AI-exposure disparities in the Pew paper was education, with over a quarter of workers with a bachelor's degree or more working in high-exposure jobs. Only about 3% of workers with less than a high school degree worked in high-exposure jobs. Across race, age, citizenship status, and even education, Pew found that average hourly earnings of high-exposure workers exceeded earnings of low-exposure workers.

SECTION 3: Anthropic data

In February 2025, Anthropic published findings on worker exposure to artificial intelligence to occupation as a part of their "Economic Index" and corresponding paper. By analyzing trends in conversations with their chatbot, Claude.ai, they measure the frequency at which certain tasks are exposed, augmented, or automated by AI use. Like the Pew Center, they rely on the O*NET work activity framework to sort their findings into occupation-level information. Unlike Pew, they begin at the more granular task level.

Anthropic links queries made to Claude.ai to one of 19,530 O*NET tasks. They measure exposure by task as the proportion of user conversations with Claude.ai that concern a certain task. For example, 2.68% of all requests made to Claude.ai during the survey period matched the O*NET task "modify existing software to correct errors, allow it to adapt to new hardware, or to improve its performance". As a result, this task is assigned a task exposure score of 2.68%. Specific behaviors which measure either automative or augmentative exposure are additionally linked to 3364 tasks.

By grouping these tasks to their associated O*NET occupations and connecting those to worker demographics, we produce a quantitatively-backed update to Pew's original assessment. All 19,530 tasks can be sorted into 37 out of the 41 Work Activities (O*NET does not assign tasks to four work activities). We group each task by its associated work activity and find the total percentage of Claude.ai queries that match each activity. Table 1 lists the top and bottom five most exposed work activities to artificial intelligence and compares these findings with Pew's initial predictions.

Work activities are exposed to artificial intelligence at different levels

Top and bottom five work activities most associated with Claude.ai use, 2024

Rank	Work activity	Exposure to AI automation	Exposure to AI augmentation	Total exposure to AI	Pew exposure estimation (2023)
1	Thinking creatively	7.2%	7.9%	15.0%	High
2	Documenting/recording information	5.0%	6.6%	11.6%	High
3	Providing consultation and advice to others	2.8%	7.5%	10.3%	Medium
4	Analyzing data or information	2.2%	5.5%	7.7%	High
5	Judging the qualities of objects, services, or people	3.2%	3.1%	6.2%	Medium
33	Identifying objects, actions, and events	0.0%	0.1%	0.1%	Medium
34	Establishing and maintaining interpersonal relationships	0.0%	0.1%	0.1%	Low
35	Staffing organizational units	0.0%	0.0%	0.1%	Medium
36	Scheduling work and activities	0.0%	0.0%	0.0%	High
37	Operating vehicles, mechanized devices, or equipment	0.0%	0.0%	0.0%	High

Source: Anthropic, O*NET, and authors' calculations
Note: Exposure is the proportion of queries made to Claude that can be associated with that work activity

Equitable Growth

These findings indicate that the Pew Center's 2023 approximation of work activity AI exposure is not wholly consistent with what our calculations using Anthropic's measurements reveal in 2025. In particular, some work activities that the Pew Center expected to be highly exposed to AI do not frequently appear in conversations with Claude.ai. The most dramatic difference is for "Operating vehicles, mechanized devices, or equipment". While Pew anticipated that artificial intelligence would be highly impactful on this component of work, it is not currently associated with any queries made to Claude.ai. Consequently, this work activity is, at the very least, not exposed to the artificial intelligence made readily available by chatbots.

Pew's estimates are more closely aligned to the higher end of our ranking. The most exposed work activity, "Thinking creatively", which is linked to 15% of all conversations in the sample, is also diagnosed as "highly exposed" by the Pew Center. Likewise, the other most exposed work activities by our measurements are either highly or moderately exposed according to the Pew Center's standards.

Despite the fact that Anthropic data reveals exposure by work activity that deviates to some degree from Pew's 2023 expectations, we find that AI exposure by race, gender, and education is much more closely aligned. Pew extrapolates its work activity levels to related Census data. We likewise compile Anthropic task data to the occupational level, then link these findings to census occupational metrics.

SECTION 4: Our method

While Anthropic provides its summed AI exposure numbers by occupation, we elect to calculate our own findings using their publicly available task exposure metrics. Like Anthropic, we first calculate exposure by occupation, using the list of occupations provided by O*NET version 20.1 from 2015. Unlike Anthropic, we also source task frequency data from O*NET, then weight each task by how often it's performed on average in its corresponding occupation. For example, the task "Interact with clients to assist them in gaining insight, defining goals, and planning actions…" occurs several times daily for Clinical Psychologists. Its exposure therefore has a stronger impact on Clinical Psychologists' total occupational exposure than the task "Provide psychological or administrative services and advice to private firms and community agencies…", which they perform yearly on average. While other studies such as the one recently published by the Yale Budget Lab or Eloundou et al. (2024) use chatbot data and the O*NET task framework to estimate worker exposure, our method of weighting by task importance when aggregating to the occupation level is –to our knowledge– unique to this study.

Once we determine overall exposure for the 20.1 O*NET occupations, we match this list to their most recent set of vocations, O*NET 30.0. From there, we crosswalk the data to the 2024 Current Population Survey. AI exposure metrics on O*NET occupations are matched and aggregated into Census occupations (Table 2).

Assessing AI exposure by Census occupation is advantageous due to the abundance of federal data on these professions. At this level, exposure can be linked to occupation, industry and worker demographics. Indeed, our assessment of AI exposure by worker demographics, and the following analysis on AI and worker income, is made possible through this linking. But there are some noteworthy features of this aggregation. The 708 O*NET occupations for which we have exposure data are assigned to only 415 available Census occupations for 2024. Consequently, some individual O*NET exposures are coalesced into a single Census exposure. For most, this is a matter of summing two or three similar occupations, but for a few, the aggregation is substantial. In the O*NET data, there are 37 vocations that fall under the CPS category of postsecondary teacher (e.g. Business teachers, Engineering teachers, Geography teachers, etc.). Therefore, while 20% of Claude.ai queries are associated with postsecondary teaching, it would be inaccurate to say that every kind of postsecondary teacher is 20% exposed. Instead, we might say that postsecondary teachers will shoulder 20% of the AI exposure facing the US labor market, if Claude.ai is truly an indicator of economy-wide uses of AI. At the level of O*NET occupations, postsecondary teachers individually make up anywhere from 0.3% of queries (Computer science teachers, postsecondary) to 1.2% (English language and literature teachers, postsecondary) of Claude.ai conversations.

Census occupations are exposed to AI use at different rates

Percentage of queries made to Claude.ai that are linked to each occupation or occupational group

Most exposed occupations	Automative exposure	Augmentative exposure	Total exposure	Least exposed occupations	Automative exposure	Augmentative exposure	Total exposure
Postsecondary teachers	7.5%	12.5%	20.0%	Hazardous materials removal workers	0.000%	0.003%	0.003%
Software developers	2.6%	4.6%	7.2%	Millwrights	0.000%	0.003%	0.003%
Computer programmers	2.8%	3.9%	6.7%	Dispatchers, except police, fire, and ambulance	0.000%	0.003%	0.003%
Web developers	1.3%	1.8%	3.1%	Crossing guards and flaggers	0.003%	0.000%	0.003%
Network and computer systems administrators	0.8%	1.6%	2.4%	School bus monitors	0.003%	0.000%	0.003%
Computer systems analysts	0.8%	1.6%	2.4%	Telecommunications line installers and repairers	0.000%	0.003%	0.003%
Computer occupations, all other	0.7%	1.2%	1.9%	Emergency management directors	0.003%	0.000%	0.003%
Tutors	0.9%	0.9%	1.8%	Railroad conductors and yardmasters	0.000%	0.003%	0.003%
Artists and related workers	0.7%	0.8%	1.4%	Laundry and dry-cleaning workers	0.000%	0.002%	0.002%
Statistical assistants	0.6%	0.8%	1.4%	Landscaping and groundskeeping workers	0.000%	0.002%	0.002%

Source: Anthropic, O*NET, and authors' calculations

Note: The "Postsecondary teachers" exposure is the sum of the exposures of 37 individual post-secondary teaching positions.

Equitable Growth

SECTION 5: Summary statistics

We used the newly weighted Anthropic data to produce figures comparable to those in the original Pew paper, confirming greater AI exposure among women, Asian, and highly educated workers.

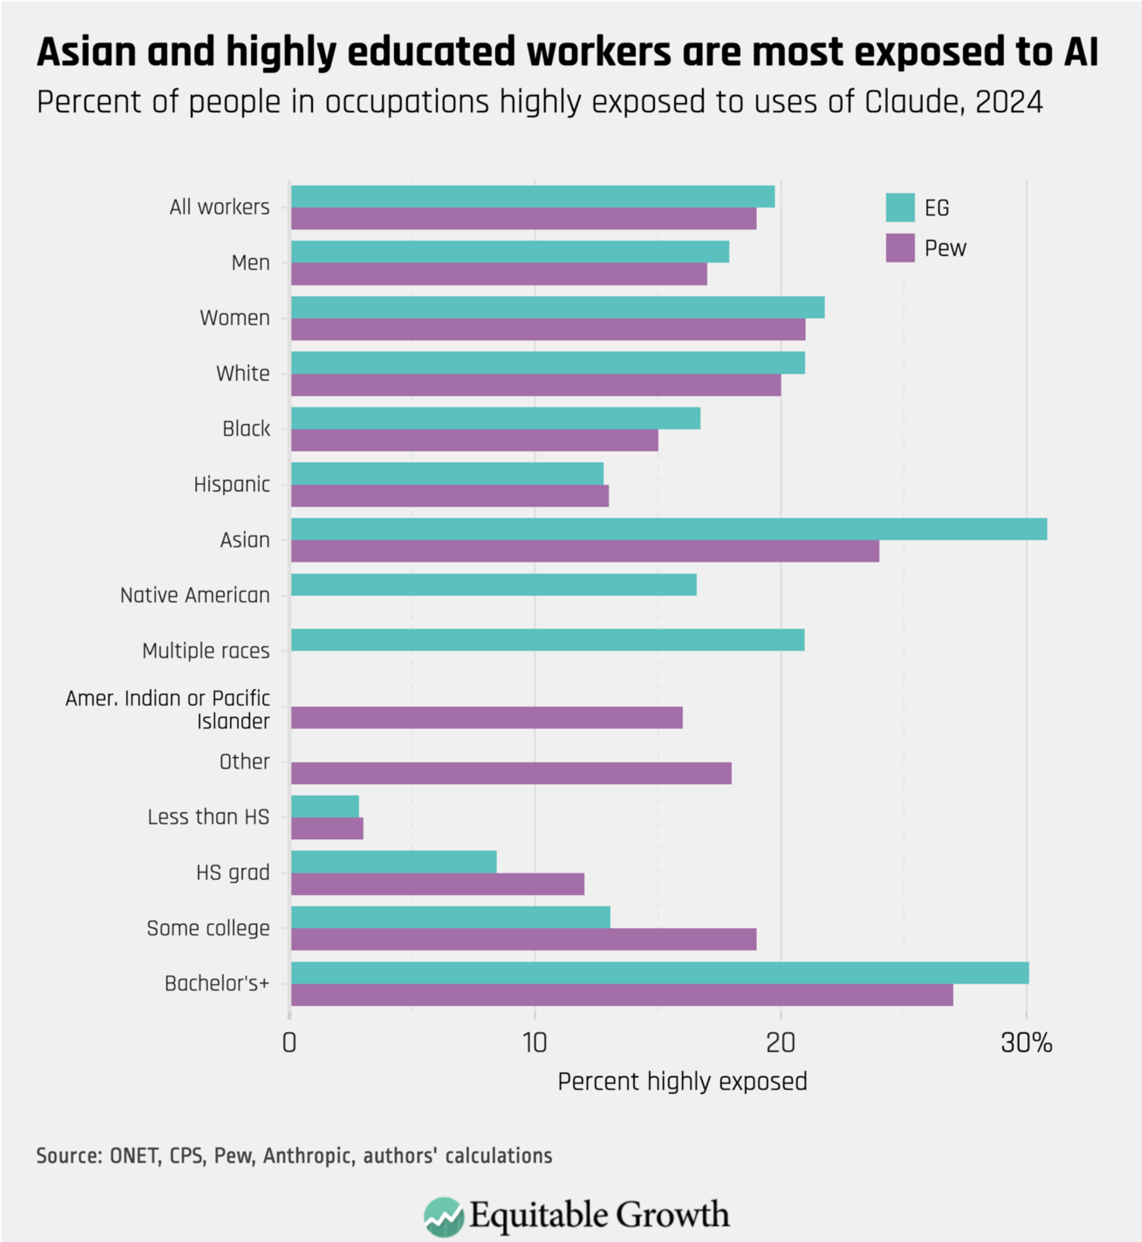

Our results show a much more significant disparity in AI exposure between Asian workers and the average worker compared to the Pew study. We used a different race variable from the Pew study, so some of the effect captured in Pew's "American Indian or Pacific Islander" category could be captured in our "Asian" category. Our results also show a larger disparity by education, with high school graduates and workers with some college education less exposed than in the Pew study. 30% of workers with at least a bachelor's degree are employed in high-exposure jobs in our analysis, compared to 27% in the Pew study. Similarly, we confirm the Pew finding that workers with higher incomes tend to be more exposed to AI, regardless of race, age, citizenship status, and education.

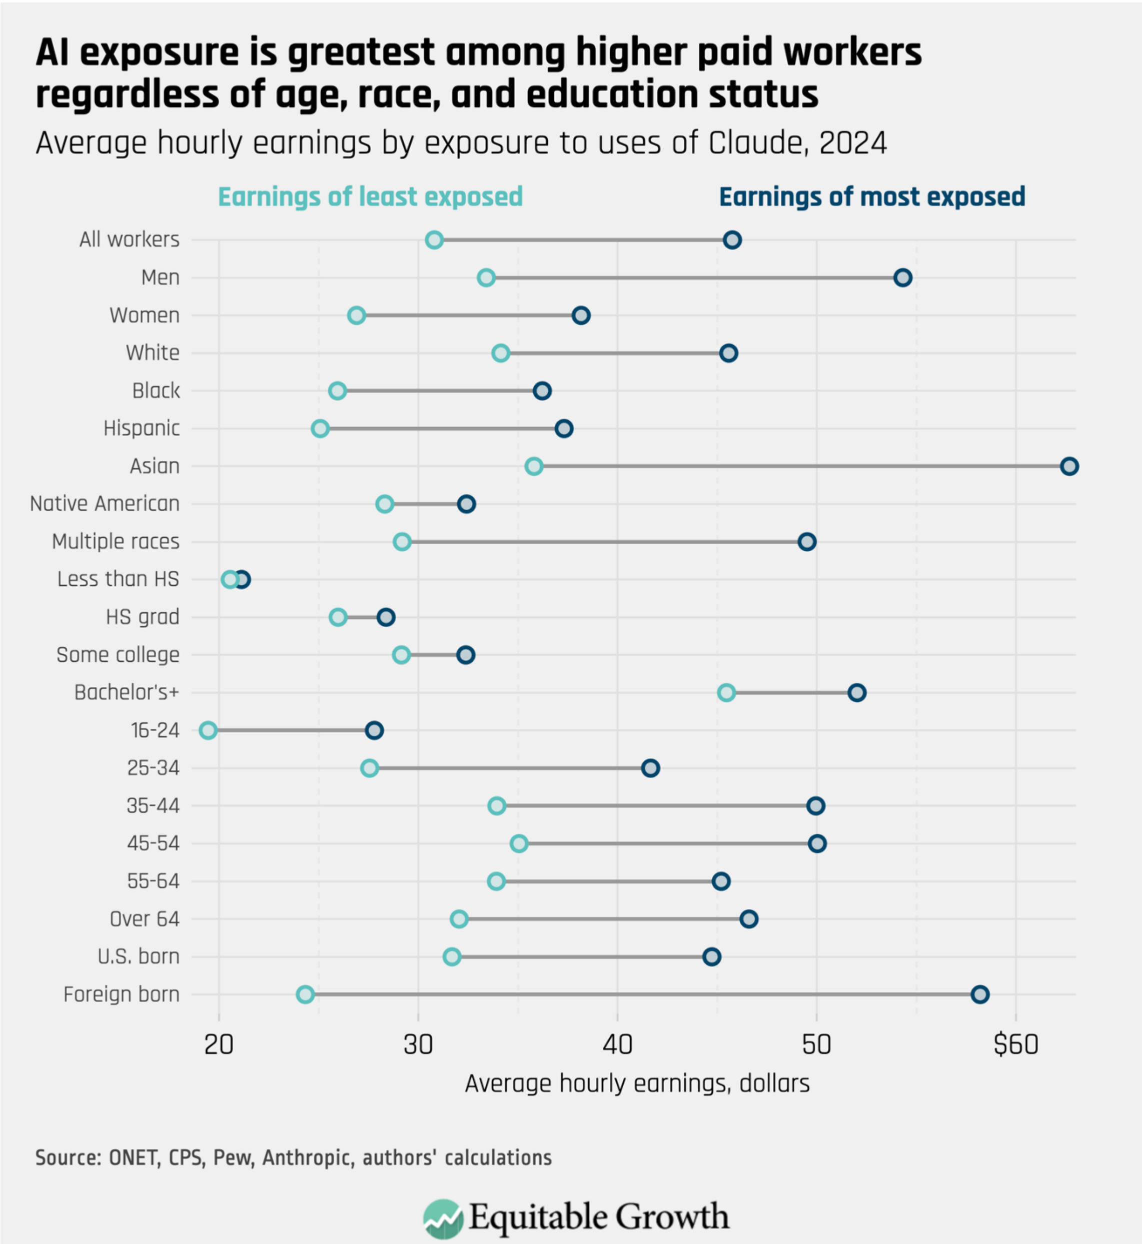

Earnings as a whole are higher in our analysis than in Pew's, which is simply a reflection of the elevated rate of nominal wage growth over the last few years. Our findings largely corroborate Pew's, with earnings disparities between most and least exposed workers greatest among Asian, foreign born, and people ages 35-44. Earnings disparities by exposure appear to increase alongside education, meaning that workers with at least a bachelor's degree could face the biggest nominal hit to income as a result of AI exposure. In other words, workers displaced by AI exposure—likely automative—will move from occupations with a higher average wage into less exposed and lower-paying occupations. Given the large nominal income gap for highly educated workers, that cohort is likely to experience the largest nominal loss from displacement. Overall, the consistent earnings disparity suggests that job loss from AI automation could push labor income down as workers move from more lucrative and highly exposed jobs to less lucrative and relatively unexposed jobs.

In addition to reproducing some of the Pew figures with the weighted Anthropic data, we produced a series of new figures that further flesh out the incidence of AI exposure along dimensions of race, gender, and education. First, we find that the direction of gender disparities in exposure is largely consistent across racial groups, with women more exposed than men in all groups except Asian workers.

Gender disparities in AI exposure are significant among White, Hispanic, and Asian populations

Percent highly exposed to uses of Claude by race and gender, 2024

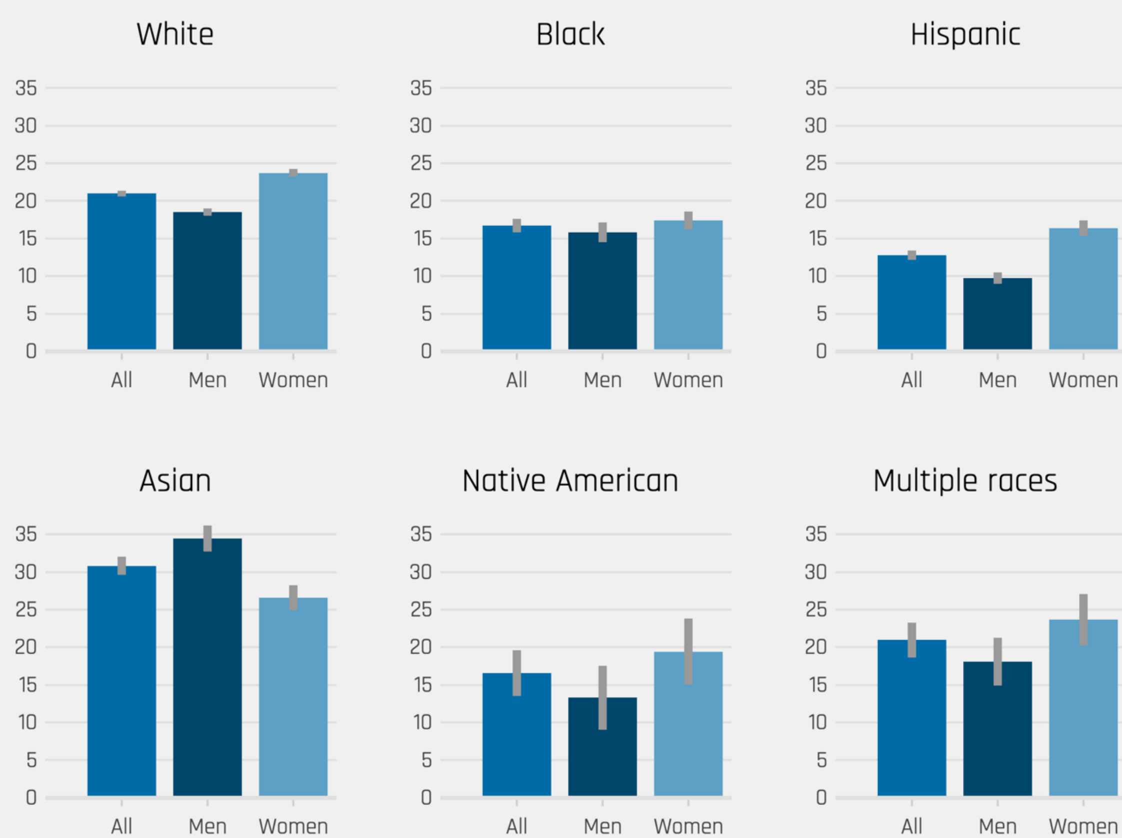

Source: ONET, CPS, Pew, Anthropic, authors' calculations
Note: Gray bars represent 95% confidence intervals

Equitable Growth

Using 2024 CPS data, gender differences are not statistically significant at the 95% confidence level for Black, Native American, and workers of multiple races. The inverted gender disparity among Asian workers could be partially explained by the skew in educational attainment among Asian workers towards men, in comparison with a national average bias towards women. This is important because gender differences in AI exposure become negligible among highly educated workers, so a male-skewed cohort of highly educated workers would tend to push down women's average exposure within the group as a whole. In other words, gender differences by educational attainment can warp measures of exposure by gender. This also suggests education could be a more direct driver of exposure disparities.

Gender disparities in AI exposure shrink as education increases

Percent highly exposed to uses of Claude by gender and education, 2024

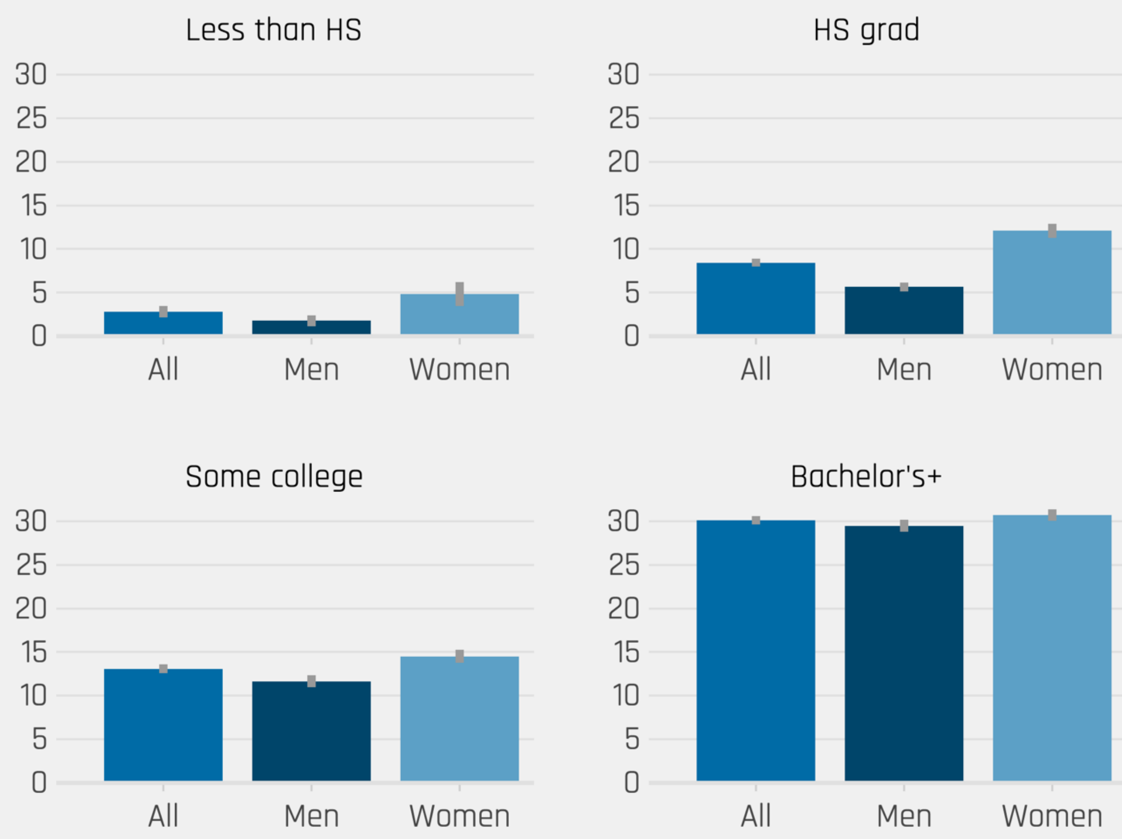

Source: ONET, CPS, Pew, Anthropic, authors' calculations
Note: Gray bars represent 95% confidence intervals

Equitable Growth

Looking at education disparities within racial groups provides further evidence of education as a fundamental determinant of AI exposure. Even within the Native American and Multiple Races categories, for which small sample sizes make the exposure-education relationship less obvious, the difference between high school graduates and holders of advanced degrees is stark. Among White, Black, Hispanic, and Asian workers, the difference in AI exposure is large and statistically meaningful across all education categories.

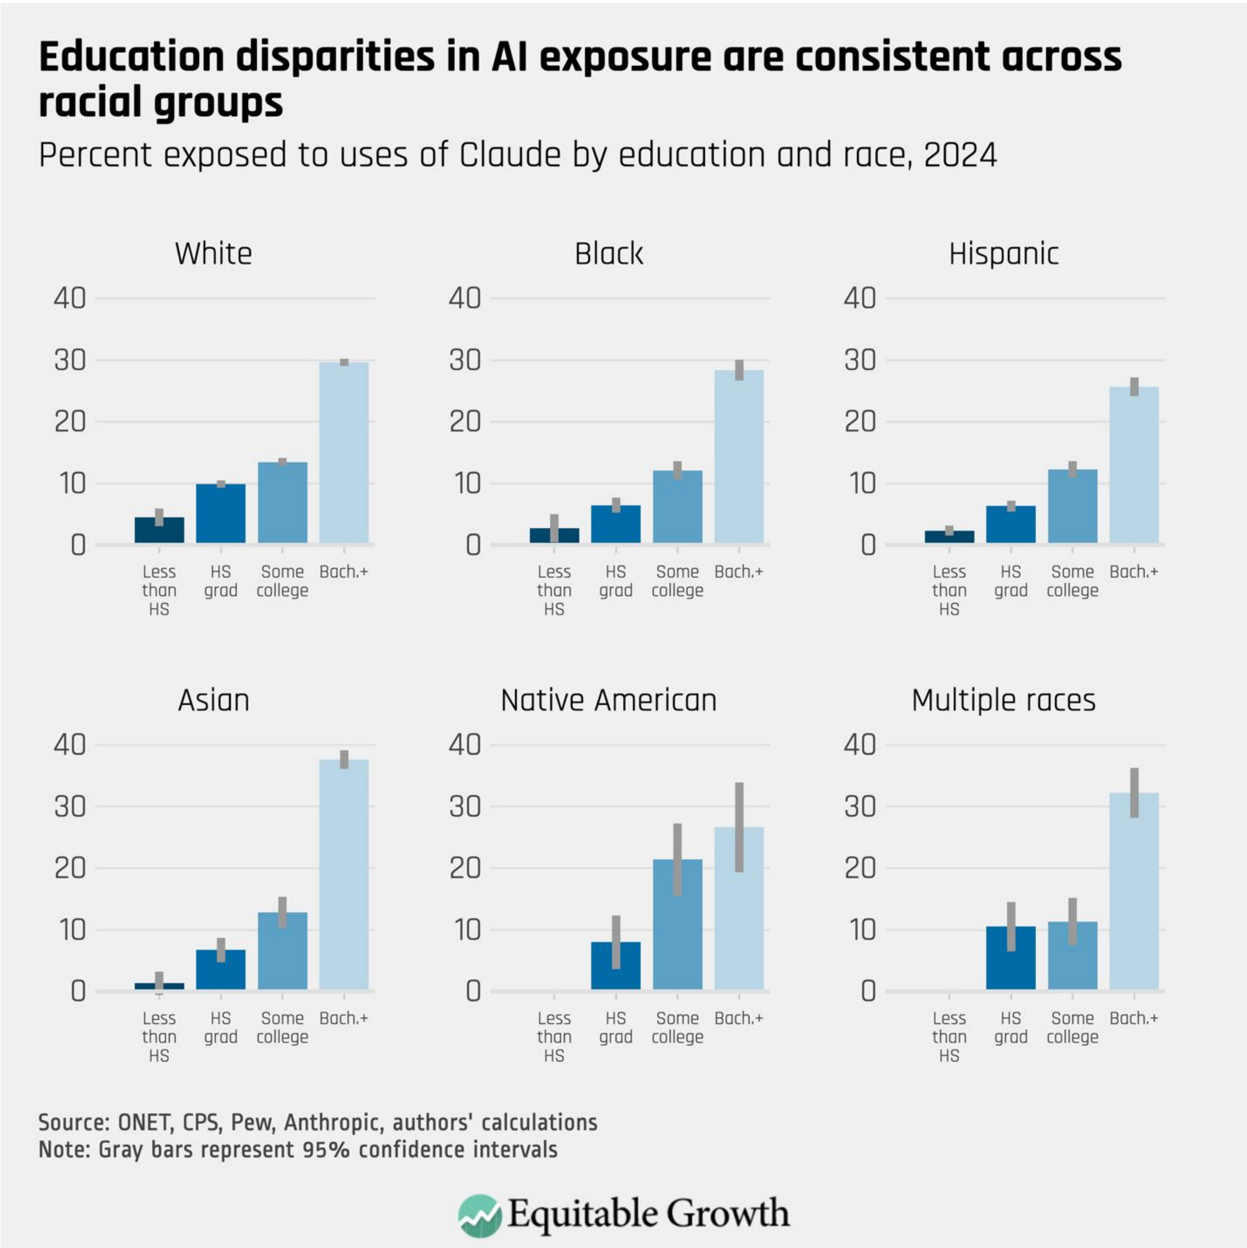

SECTION 6: Regression

The data reveal that exposure to artificial intelligence can vary by age, occupation, education level, and gender. But whether these variations in exposure meaningfully contribute to differences in income is still largely an unanswered question. In general, is higher AI exposure associated with higher or lower wages, and how does this relationship change by type of exposure?

Our analysis reveals that currently, exposure to artificial intelligences is not a substantial determinant of workers' wages. But the story changes by exposure type. Automative and augmentative AI exposures have stronger, but counteracting effects on wages. Whether artificial intelligence improves or undermines worker outcomes is largely determined by whether AI supports or replaces worker tasks.

Figure 6 reveals the impact of exposure to artificial intelligence on wages. Here, the percent change of hourly wages is calculated per a percent change in exposure (log-log regression). Model one (seen in blue) uses total exposure, while model two (seen in orange) differentiates automative exposure from augmentative exposure. Controls for O*NET job content, demographic, and economic characteristics are added, and the regression is weighted by demographic populations within occupations.

Overall exposure to AI does not have a large impact on wages, but automation and augmentation AI exposure have greater, opposing wage effects

Percent change in hourly wages given a percent change in AI exposure by type

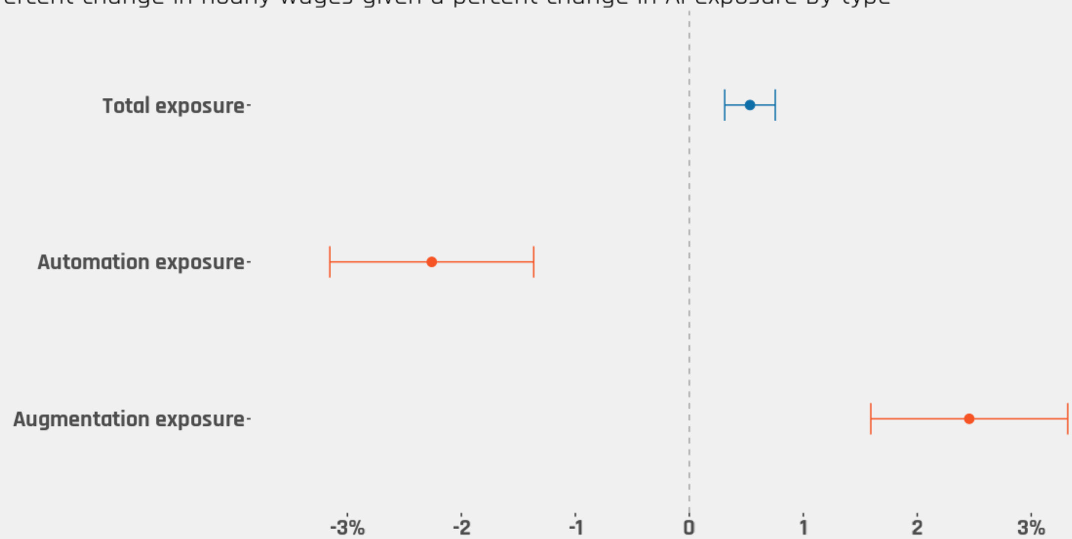

Source: Authors' calculations using Anthropic, Current Population Survey and O*NET data.

Note: Estimates are coefficients from a regression of log real wages on AI exposure with controls for O*NET job content importance and individual demographic and economic characteristics, expressed in percent terms. The regression is weighted by population demographics within occupations. The coefficient on model one is shown in blue, the coefficients on model two are shown in orange. The bars represent 95% confidence intervals

Equitable Growth

All else held equal, a percent increase in total AI exposure will on average predict an 0.5% increase in hourly wages. This finding is additionally consistent with Figure 2, which shows that workers who are more exposed to AI also tend to work in higher paid professions.

The current impact of total exposure to AI on wages, though statistically significant, is relatively small. It has less influence on wages than industry, age, education, gender, race and job content. However, what is notable is that the relatively minimal overall effect of AI exposure on wages may conceal the larger effects of individual exposure types. Our second model demonstrates that, when considered on its own, exposure to automative AI reduces wages by 2.3%, but this effect is simultaneously counteracted by impact of augmentative AI exposure, which increases wages by 2.5% per percent increase in exposure, all else held equal.

SECTION 7: Discussion

These results are only a small picture of the state of the labor market during the advent of AI. They are derived from queries to only a single chatbot in a market rapidly flooding with various forms of artificial intelligence. But the results are telling. AI has the potential to improve some workers' wage outcomes if it helps them do their jobs. But it is equally likely to reduce worker wages in areas where it substitutes them. Academics and policymakers alike may therefore find it helpful to study not only where AI is entering the workforce, but how.

As indicated in Table 2, the CPS occupations that are highly exposed to automative AI are also more likely to be highly exposed to augmentative AI. However, the data also indicate that occupations whose tasks are more likely to be augmented than automated also tend to have higher hourly wage rates. Figure 7 demonstrates this phenomenon visually. Here, occupations' augmentative AI exposure scores are graphed against their automative scores. The majority of occupations have low exposure rates that stick closely to the 1:1 line (where automative exposure equals augmentative exposure). But there are a number of points that are distinctly below the line of parity, indicating that augmentative exposure far surpasses automative exposure for their corresponding professions. Notably, the majority of these points (and most profound examples) all represent occupations whose hourly wage exceeds the median rate of our sample ($34.25). To better show this phenomenon, outlying data points representing Postsecondary Teachers (total exposure = 20.0%, average wage = $51 per hour), Software developers (total exposure = 7.2%, hourly wage = $72), and Computer Programmers (total exposure = 6.7%, average wage = $48) are all excluded from the figure. Each of these occupations are 40-74% more exposed to AI augmentation than automation and lie below the equilibrium line.

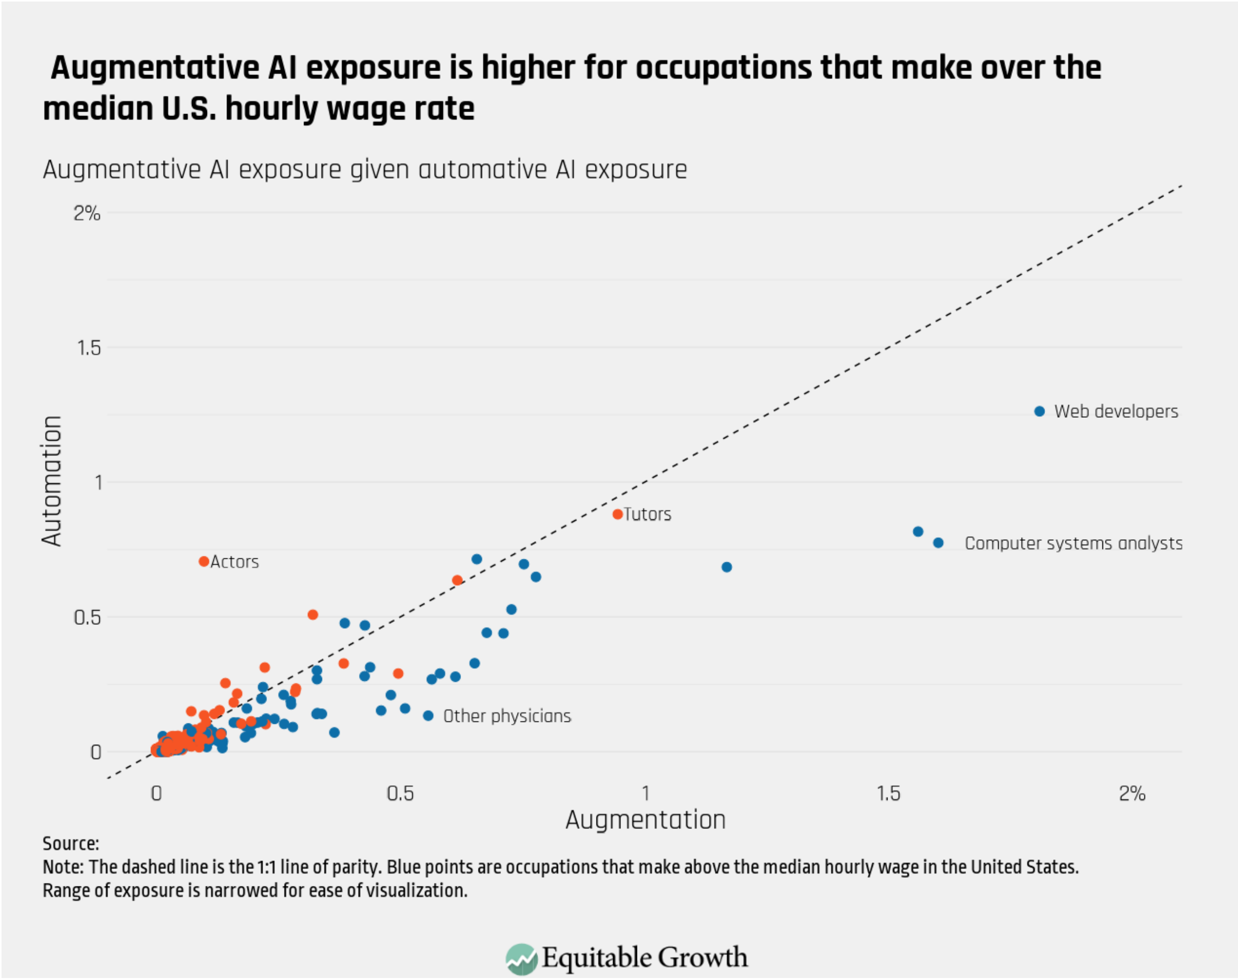

The wage-inflating effect of augmentative AI exposure over automative AI can also be seen in tables 3 and 4. Table 3 lists the occupations where augmentative exceeds automation by the highest amount. Many of these vocations (Postsecondary teachers, Software developers, Web developers, etc.) also have the highest total exposure out of the CPS occupation list. The average pay per hour of these ten occupations is $55 per hour, which greatly exceeds the overall mean wage of $37.62.

Occupations where augmentative exposure greatly exceeds automative exposure tend to pay more than other occupations Top 10 occupations by amount augmentative exposure exceeds automative				
Occupation	Automative exposure	Augmentative exposure	Total exposure	Average pay per hour
Postsecondary teachers	7.54%	12.49%	20.03%	$51.16
Software developers	2.64%	4.58%	7.22%	$71.68
Computer programmers	2.79%	3.92%	6.71%	$47.96
Computer systems analysts	0.77%	1.60%	2.38%	$49.38
Network and computer systems administrators	0.82%	1.56%	2.38%	$43.55
Web developers	1.26%	1.81%	3.07%	$50.22
Computer occupations, all other	0.68%	1.17%	1.85%	$47.49
Other physicians	0.13%	0.56%	0.70%	$95.95
Biological scientists	0.16%	0.51%	0.67%	$42.12
Web and digital interface designers	0.28%	0.61%	0.89%	$48.10

Source: Anthropic, O*NET, and authors' calculations

Note: The "Postsecondary teachers" exposure is the sum of the exposures of 37 individual post-secondary teaching positions.

Equitable Growth

Conversely, the occupations in which automative exposure exceeds augmentative (Table 4) exposure tend to pay less than average. The mean wages per hour of these ten occupations is $33 per hour.

Occupations where automative exposure exceeds augmentative exposure tend to pay less than other occupations Top 10 occupations by amount automative exposure exceeds augmentative				
Occupation	Automative exposure	Augmentative exposure	Total exposure	Average pay per hour
Actors	0.71%	0.10%	0.80%	$33.38
Computer numerically controlled tool operators and programmers	0.51%	0.32%	0.83%	$26.55
Correctional officers and jailers	0.25%	0.14%	0.40%	$27.59
Archivists, curators, and museum technicians	0.48%	0.39%	0.86%	$41.02
Special education teachers	0.31%	0.22%	0.53%	$32.95
Other woodworkers	0.15%	0.07%	0.23%	$23.07
Software quality assurance analysts and testers	0.71%	0.66%	1.37%	$44.74
Word processors and typists	0.22%	0.17%	0.38%	$24.70
Legal support workers, all other	0.06%	0.01%	0.07%	$42.59
Elementary and middle school teachers	0.47%	0.43%	0.90%	$36.13

Source: Anthropic, O*NET, and authors' calculations

Equitable Growth

SECTION 8: Conclusion

Worker exposure to artificial intelligence, measured here in terms of uses of Anthropic's Claude.ai agent, has an ambiguous relationship with income. An occupation's exposure to augmentative Claude.ai use is highly positively correlated with income, while exposure to automative use is highly negatively correlated. Overall exposure is slightly positively correlated with earnings, due in part to Claude.ai uses tending to be labeled augmentative.

All else held equal, a percent increase in total exposure on average predicts a 0.5% increase in hourly wages.

Corroborating earlier work from the Pew Research Center, we find meaningful differences in AI exposure along lines of gender, race, education, and income. Women tend to work more in highly exposed occupations compared to men, a relationship that holds across most racial groups. Workers with college educations or high incomes also tend to work in highly exposed occupations, trends that are persistent across race. White and Asian workers overall tend to work more in highly exposed occupations compared to Black and Hispanic workers.

Our work advances the growing body of research on labor market impacts of artificial intelligence, and suggests avenues for future exploration. Different types of AI use are likely to be deployed to differing effect in the labor market, with some uses assisting human workers and other uses displacing human work altogether. Policymakers should be cognizant of these differences when regulating the development and workplace deployment of artificial intelligence.