

Yale Budget Lab finds no significant AI effect on US jobs or wages through Q1 2026

Source copy kept by wheresthe.ai for the story "Yale Budget Lab finds no significant AI effect on US jobs or wages through Q1 2026" (<https://wheresthe.ai/d/us-yale-budget-lab-no-ai-labor-effect-2026-05/>).

Source: <https://budgetlab.yale.edu/research/what-we-do-and-dont-know-about-how-ai-affecting-labor-market>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 0f3bf521383e64247138c84bdc4ae02280543dddf4fd84196740552756a21ab7

Key Takeaways

AI-exposed and unexposed occupations are quite different, making it challenging to learn about AI effects by comparing them

One such difference is that AI-exposed occupations are generally less sensitive to the business cycle than unexposed occupations

We use a synthetic differences-in-differences design that addresses these challenges

So far, this approach provides no clear evidence of AI effects on the labor market, but this could change quickly

The rapid adoption of AI tools into the workplace has been met with a mixture of excitement and anxiety, with much of the latter focusing on potential labor market disruption. Prior Budget Lab analysis and ongoing tracking efforts have highlighted labor market outcomes, like occupational churn, for which one could reasonably expect new AI tools to leave a footprint in the labor market. But Budget Lab efforts have not yet included estimates of the effects that AI tools are having on the labor market. Doing so is complicated by the fact that the macroeconomy is subject to many simultaneous changes and that any effects of AI are likely operating on a slice of the labor force that is quite different than the rest. Using an approach called synthetic differences-in-differences (SDID) that addresses these challenges, we generally find no statistically or economically significant effects as of yet. Specifically, we examine the employment and wages of AI-exposed and unexposed occupations, which are informative about the kinds of labor market shocks that AI might be generating. For reasons discussed below, occupation-level AI exposure metrics are by no means ideal for distinguishing the portions of the labor market that are more or less likely to experience an AI-induced negative demand shock (Ischenko and Millet 2026 ; Richmond 2026). AI exposure metrics are not necessarily indicative of a job at risk of automation, but rather a job whose tasks are especially likely to be affected by AI. This means that analysis of exposure metrics could actually hide effects if (for instance) half of "exposed" occupations see their jobs expand and half see their jobs contract—the average could be zero. Moreover, exposure metrics are based on judgment and could exclude tasks and jobs that turn out to be affected importantly by AI. This is why Budget Lab's AI tracker uses measures of change in occupational distribution, rather than comparisons of exposed and unexposed jobs, as our primary metric for possible AI disruption. That being said, the exposure metrics are appealing to many researchers, in part because they allow for cross-sectional comparisons that help us learn about potential AI effects. To the extent that these comparisons are conducted, it is important that they accommodate persistent differences between the groups, like the higher educational attainment of workers in exposed occupations and the heightened business cycle sensitivity of unexposed occupations in historical data. As part of a larger suite of methods, such an approach can be illuminating about AI impacts. Though we find no effects as of yet, The Budget Lab will periodically update the assessment as new labor market data become available.

If we want to understand AI impacts, why isn't it sufficient to simply compare exposed and unexposed workers' outcomes today? The reason is that these occupations (and their workers) are different in ways that are unrelated, or only incidentally related, to AI exposure. According to the standard metrics, which Budget Lab has analyzed in detail previously, AI exposure is distributed unequally across workers. ¹ New large-language models are very capable of assisting with coding activities, for example, but generally cannot intervene in the physical world and have limited utility for tasks requiring such intervention. Distinguishing different types of jobs along this dimension is in principle useful. However, our AI exposure metric and others currently in use have some important limitations. These measures tend not to be specific about whether AI is likely to replace or enhance a worker's labor, or to what extent the technology is currently being used (Massenkoff and McCrory 2026). AI exposure metrics are ultimately informed guesses as to which tasks and occupations are most likely to be affected by AI tools, based on

some combination of human and LLM judgment as to which tasks overlap with LLM capabilities. This can change over time and could be wrong from the start. Those caveats in mind, it is remarkable how unequally distributed AI exposure is across the labor market. Workers in exposed jobs have persistently different outcomes than their counterparts, as shown in Table 1. We consider occupations to be exposed to AI when they are in the top third of occupations by the Budget Lab's summary metric, and unexposed when they are in the bottom third. 2 As measured by number of distinct occupations, the most common major occupational group in the exposed category is Office and Administrative Support; in the unexposed category, the most common group is Protective Services. Workers in exposed occupations are more likely to be women and much more likely to have a four-year college degree. Based on pre-pandemic data, they are also less exposed to the business cycle, suffering smaller reductions in employment during recessions. Larger coefficients indicate a more procyclical pattern, with employment rising during booms and falling during downturns. (See the appendix for more discussion of this.) Exposure is also correlated with other factors which may affect hiring and firing, like remote work (Kolko 2026).

These gaps make it inherently difficult to interpret comparisons between exposed and unexposed occupations. 3 Fortunately, there exist econometric strategies that can be useful for handling at least some of these differences, whether observed or unobserved in our data. Our preferred strategy, synthetic differences-in-differences, is an amalgam of the better-known synthetic control and differences-in-differences methods. It allows one to construct an apples-to-apples comparison group for AI-exposed workers using a weighted combination of unexposed occupations. Differences between this comparison group and exposed workers can, under certain assumptions, be interpreted as effects of AI. 4 Most importantly, one must assume (as discussed above) that AI exposure metrics have some ability to distinguish, however imperfectly, occupations that are more or less likely to be affected by AI.

We now apply this strategy by showing the average employment share of an occupation (out of the civilian population) in the AI-exposed group against the average share of the synthetic comparison group. Figure 1 shows these shares over time, with the dashed vertical line indicating the fourth quarter of 2022, when ChatGPT was released for a mass audience. 5

Because SDID has the desirable feature of accommodating level differences in the AI-exposed and unexposed groups, it can be difficult to visually determine the difference in the two groups: the estimated impact of exposure to AI. Figure 2 shows that difference, along with confidence intervals that indicate whether the impact in a given quarter is significantly different from zero. 6 No impact is evident through the post-2022 time window. 7

We also examine the log real hourly wages of AI-exposed and unexposed workers. 8 Figure 3 shows these trends, and Figure 4 shows the impact of LLM introduction. Here again, we see no statistically or economically significant impact.

In appendix figures, we additionally show results for the unemployment rate, for our entire sample and for a subgroup of workers ages 16 to 34. These figures indicate positive effects in the most recent quarter—roughly half a percentage point increase in the entire sample, and more for the 16-34 year old subsample—but both are statistically insignificant as of the first quarter of 2026. It is possible that AI has raised unemployment for workers in exposed occupations; however, unemployment may not be the ideal outcome to track for this purpose. Unemployment effect estimates are especially imprecise in our context and subject to the limitation that occupation measurement is complicated for unemployed workers, many of whom are not associated with a specific occupation.

A small dose of economic theory is helpful for interpreting the results above. In the labor market, there are four basic patterns that could characterize employment and wages, each of which is associated with a different kind of labor market “shock”. The possibilities are: Employment rises and wages fall: a positive supply shock Employment falls and wages rise: a negative supply shock Employment and wages both rise: a positive demand shock Employment and wages both fall: a negative demand shock The pattern that many anticipate, for AI-exposed workers, is the last one: a negative demand shock. As AI tools allow employers to automate certain tasks and the workers who perform them, one could expect employer demand for labor to fall, and both employment and wages to decline. Within the synthetic differences-in-differences framework, we do not yet see evidence of this pattern. However, there are some important caveats that should be kept in mind. First, LLMs are constantly improving, such that the economic impacts from early models seem likely to be smaller than those of recent models. Second, the validity of the SDID approach taken in this analysis depends on whether the introduction of LLMs can be distinguished from other events, like the rise in remote work, that are closely correlated with exposure to AI. 9 It also depends on how well AI-exposure metrics proxy for the tasks and jobs actually at risk in the age of AI. If “AI-

exposed” occupations include some that receive negative demand shocks and some that receive positive shocks, our analysis would miss the former cause as it averages across the whole group. Finally, it worth noting that the Bureau of Labor Statistics’ Current Population Survey, our data source throughout this analysis, is best suited to analysis of broad groups and somewhat underpowered for subgroup analysis, like the 22–27 year old recent college graduates that have been a special focus in some AI labor market impact discussions. 10 If labor market effects of AI are currently limited to a narrow slice of the workforce, other datasets and research designs would be better suited to identify them. Using this and other econometric approaches, and keeping in mind their diverse strengths and weaknesses, the Budget Lab will continue to monitor the labor market for signs of AI effects.

The code used to perform this analysis can be found [here](#). In this appendix we first describe the synthetic differences-in-differences methodology in more detail, briefly discuss our approach to understanding occupation employment cyclicity, and then provide analysis for additional labor market outcomes not shown above.

Synthetic differences-in-differences was developed by Arkhangelsky et al. (2021). We use an R implementation described [here](#). “Treated” occupations are those in the top tercile of occupations, and the donor pool of untreated occupations are those in the bottom tercile. Prior to implementing this procedure, we adjust for seasonal variation at the detailed occupation level by purging the effects of quarter-of-the-year dummies, estimated over the entire sample. SDID requires assumptions broadly similar to those needed for any econometric method that constructs a never-treated comparison group for a treatment group that receives treatment during the sample window. In the absence of the treatment, the (unobserved, counterfactual) outcome growth path for treated occupations is assumed to be reproducible with a suitably chosen synthetic comparison group. One assumption it does not require, in our context, is that workers in exposed and unexposed occupations would all have had the same labor market outcomes, absent the introduction of new AI tools. SDID allows for persistent differences between exposed and unexposed occupations, and it finds a weighted mix of time periods and unexposed occupations that arguably makes the parallel trends assumption more plausible. Because SDID allows for level differences between the treatment and control groups, it assigns positive weights to a wider array of units within the control group than does traditional synthetic control. In our context, the variation in weights across untreated units is relatively small: for example, Building and Grounds Cleaning and Maintenance occupations have about 75% the weight given to Office and Administrative Support occupations when untreated occupations in those groups are used to create a synthetic comparison group for treated occupations’ mean employment share. Across the synthetic (i.e., weighted) comparison groups for our different outcome variables, the mean AI-exposure values are quite similar, all clustering around the 20 th percentile of AI exposure. (This difference between this value and the mean exposure value of unweighted control group occupations is negligible.) By contrast, the treated group mean exposure value is about the 85 th percentile of that metric. A strength of SDID, especially in our context, is that it generates a clear visual representation of both the pre-AI “fit” of the comparison group and a quarter-by-quarter picture of how the groups are (or are not) diverging after AI was introduced. This is particularly valuable given that AI tools have become more powerful and ubiquitous over time. Even if no effects are evident in 2022 or 2023, they may yet become evident in 2026 or 2027. SDID is certainly not infallible. If an entirely different economic event were to differentially affect AI-exposed and unexposed occupations at the same time that AI models are introduced, this effect would be confused with that of AI. Moreover, as Budget Lab has explored in prior research, this strategy is only as strong as the AI exposure metrics that underlie it. Our estimates should be understood as showing how the AI-exposed group as a whole is evolving relative to an unexposed group that was carefully selected to mimic the exposed group’s growth path prior to 2022.

To calculate the cyclicity coefficient in Table 1, we separately regress each occupation’s log employment on the CBO output gap from 1994 through 2019. Especially procyclical occupations are those whose employment tended to rise more when the output gap is high (and the economy is strong) and fall more when the output gap is low (and the economy is weak). This historical analysis should not be confused with assessments of how AI-exposed occupation employment (or other outcomes) have evolved during the most recent episode of labor market cooling, since 2022 or 2023. For example, one analysis observes that more-exposed occupations saw larger declines in job postings than unexposed occupations since 2022. The aftermath of the pandemic was an unusual macroeconomic episode, and indeed we do not find the same pattern of reduced cyclical sensitivity for exposed occupations during the last few years.

Using the same econometric approach deployed above, we additionally examine changes in unemployment rates. One complication is presented by the well-known tendency of unemployment rate differences across groups to be “multiplicative”, in the sense that the ratio of one group’s unemployment rate to another group’s is relatively constant over time. To accommodate this pattern, we take logs of occupation unemployment rates and apply the

SDID procedure below. (Subsequently, we undo the log transformation to recover percentage point impact estimates, anchoring on the exposed group's unemployment rate level.) AI-exposed unemployment has risen somewhat more than unemployment of the comparison group, though the difference is not statistically significant.

Much of the discussion of AI labor market impacts has focused on younger workers, who may be having more trouble breaking into the labor market than others (Ozimek and Goldschlag 2026 ; Brynjolfsson, Chandar, and Chen 2025). We therefore investigated effects for workers 34 and younger, finding mixed evidence of AI effects for this group. To be clear, this is a much broader group than the early career subgroup that has been the focus of AI-related research and discussion in Brynjolfsson, Chandar, and Chen (2025) and elsewhere.

The authors are grateful to Sarah Bana, Nathan Goldschlag, Maxim Massenkoff, Adam Ozimek, Ernie Tedeschi, and Lee Tucker for insightful feedback on earlier drafts.

Footnotes

1 In this analysis, we use the PCA metric created by The Budget Lab and described in the linked paper. Results are not sensitive to the choice of this exposure metric or the commonly used Eloundou et al. (2023) metric.

2 We calculate terciles on an unweighted basis and do not use sample weights in the research design. As such, our method implicitly considers every occupation as providing an equal amount of information about the potential impact of AI. See this prior Budget Lab analysis for detailed discussion of how different metrics perform and how the Budget Lab summary metric is generated.

3 The challenge is even greater than it might appear, because observable worker-level differences are not the only ways in which workers in exposed and unexposed occupations could differ. For example, it could be that AI-exposed occupations happen to have experienced stronger long-run demand growth prior to the AI era. In this case, a simple comparison of growth of exposed and unexposed employment since 2022 would actually understate effects of AI.

4 See the appendix for more details. One recent paper taking a broadly similar econometric approach to related questions is Johnston and Makridis (2026).

5 We omit 2020 throughout our analysis, given that the unusual labor market pattern of that time is likely unhelpful for understanding post-AI trends.

6 Confidence intervals are calculated following the bootstrap procedure described in Clarke et al. (2023).

7 For context, the average employment share (of the civilian population) in an exposed occupation in the latest quarter was just above 0.16%, such that the top of the 95% confidence interval for the latest quarter is equivalent to only about 5% of that average share, and the point estimate itself in that quarter is virtually zero.

8 Because the Current Population Survey earnings topcodes are fairly restrictive (and change over time), we calculate occupation-specific medians rather than averages. Figure 3 therefore shows changes in averages of occupation median real hourly wages.

9 See Tucker (2026) for discussion of how this can make differences-in-differences strategies difficult to use for understanding AI effects.

10 In analysis not reported here, we restricted the sample to 22–27 year olds (of any educational attainment) and found that relatively few occupations met our sample size criterion of 50 observations per quarter.