AI Utilization and Economic Performance

Authors: Christos A. Makridis, Tina Highfill, and Jon D. Samuels[1]

Author of Contact: Christos A. Makridis

Date: April 2026

Abstract: We study whether AI use is associated with changes in output and labor market outcomes across U.S. states and industries. Our analysis combines two 2025 measures of AI exposure: a state by industry measure based on worker reported frequent AI use from the Gallup Workforce Panel, and an industry level measure based on firm reported AI use in production from the Census Bureau's Business Trends and Outlook Survey (BTOS). Together with annual data from the Longitudinal Employer Household Dynamics (LEHD) and Bureau of Economic Analysis (BEA), we exploit within-industry-state event study models, finding that higher worker reported AI utilization is associated with increases in employment and, even more so, real output. Using the industry variation only does not produce the same output pattern. Using a rough, aggregate approximation, the increase in the share of frequent AI users from 12 percent in Q2:2024 to 26 percent in Q4:2025 would correspond to roughly 1.4 to 2.8 percent higher real output, or about 0.9 to 1.8 percentage points of annualized growth over that diffusion window. These results are consistent with AI having positive economic effects and underscore the importance of using granular AI adoption rates.

Keywords: Artificial intelligence, technology adoption, productivity, employment, wages, output, worker surveys, firm surveys, event study

JEL Codes: O33, J23, J24, D22, E24, C23



The views expressed in this paper are those of the authors and do not necessarily represent the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce.

[1]Christos: Arizona State University, Gallup, Stanford University, cmakridi@asu.edu. Tina and Jon: Bureau of Economic Analysis. Thank you to Erik Brynjolfsson and Emin M. Dinlersoz for comments on the draft.

1. Introduction

Artificial intelligence (AI) has rapidly moved from a specialized research domain into a widely diffused general purpose technology (GPT) Brynjolfsson et al. [2021], Bick et al. [2025], Makridis [2025]. Yet its macroeconomic consequences remain uncertain and widely debated Acemoglu [2024], Aghion and Bunel [2024]. Public attention to AI accelerated sharply after 2022, but the relevant economic adjustments may have begun earlier, as enterprise tools and large language model applications were already entering production workflows by 2021 Johnston and Makridis [2025]. Unlike traditional capital investments, such as machinery or information technology equipment, many AI capabilities are not separately identifiable in most economic data because they are embedded in software, cloud services, bundled digital products, and organizational routines Highfill et al. [2025]. As a result, the contribution of AI to output, labor demand, and productivity is difficult to observe directly in conventional national accounts statistics Makridis and Brynjolfsson [2026]. This is not just technical challenge: it shapes how quickly economists can distinguish genuine macroeconomic effects from noise, timing assumptions, or gaps in official data.

Recent work has begun to address this problem from two complementary directions. Highfill and Samuels [2026a] develop early estimates of AI's economic impact using Census based industry measures of AI use linked to Bureau of Economic Analysis (BEA) and Bureau of Labor Statistics (BLS) production accounts. Their baseline results suggest that AI is productivity enhancing and input saving, although they also emphasize that alternative timing assumptions yield weaker estimates and that progress in economic measurement remains essential for drawing stronger conclusions. Johnston and Makridis [2025] trace out the economic effects of AI exposure using state by industry variation, finding that more exposed sectors experienced higher output, employment, and wage growth beginning in 2021, with the employment effects differing sharply depending on whether AI augments workers or automates their tasks. Consistent with the view that diffusion has been rapid even if aggregate labor market adjustments remain incomplete, Bick et al. [2025] document substantial workplace and household adoption of generative AI in nationally representative survey data. These papers complement an emerging empirical microeconomics literature examining the labor market effects of occupational exposure to AI using individual-level data [Brynjolfsson et al., 2025, Humlum and Vestergaard, 2025], which tends to find modest average effects alongside meaningful heterogeneity, including some adverse effects for early-career workers.

The primary contribution of this paper is to provide an empirical assessment of the association between actual AI utilization (not just exposure) on measured economic outcomes. First, we introduce the Gallup Workforce Panel. Beginning in Q2:2023, Gallup began measuring the frequency of AI utilization on a nationally representative sample of individuals. The survey continued annually in Q2:2024, but in Q2:2025, the survey began asking about AI utilization quarterly. When we benchmark the Workforce Panel against the Census Bureau's Business Trends and Outlook Survey (BTOS) that asks firms about their AI use [Bonney et al., 2024], we find that the BTOS may be understating the share of AI use by roughly 15 to 25 percentage points when compared to the broadest measure of AI use in the Gallup

data, although after the BTOS AI question was revised with broader language in late 2025 it more closely matches the share of frequent AI users in the Gallup data. Understanding the extent to which industries and workers are using new technologies like AI is essential for correctly measuring important macroeconomic outcomes like productivity.[2] If surveys based on firm responses consistently display lower AI usage rates than surveys based on employee responses, then analyses using firm-based surveys of AI use may understate the measured associations between AI adoption and economic outcomes. Since many of the official economic surveys conducted by statistical agencies like the Census Bureau are sent to businesses and not directly to employees, this finding could have broader implications surrounding the use of business surveys to detect adoption of new technologies. Importantly, our empirical strategy to estimate the impact of AI relies on measured differences in AI use across industries, not the level of AI use within industries. Thus, differences between the Gallup Panel and the BTOS would only impact our results inssofar as relative industry-level use differs in the two surveys.

Second, we trace out the evolution of employment, real output, and productivity. Our identification strategy exploits variation within the same state and industry over time, controlling for all state×industry and industry×year shocks. We find that a 1 percentage point increase in the AI share of users in a state×industry is associated with a 0.2-0.4% and 0.1-0.2% increase in employment and real output, respectively, following the adoption of ChatGPT. These estimates are broadly in line with the results from Johnston and Makridis [2025] using exposure to AI, although as time elapses, exposure and AI use are likely to deviate. These magnitudes are economically consequential, but not implausible, once mapped onto the observed pace of diffusion in AI use. For example, an increase in the share of frequent AI users from 12 percent to 26 percent would imply roughly 1.4 to 2.8 percent higher real output under our estimates.[3] Annualized over that diffusion window, this back of the envelope calculation implies about 0.9 to 1.8 percentage points of real output growth. Relative to recent macroeconomic benchmarks, this places our reduced form estimates above the more conservative aggregate calibration in Acemoglu [2024], who projects only modest GDP gains over a ten year horizon, and closer to the more expansionary view in Aghion and Bunel [2024], although the comparison is not exact because our estimates are reduced form state by industry associations rather than an aggregate TFP calibration.

[2] Although how much measurement strategies affect inference will depend on the identification, in our setting, differences in AI usage rates will influence the relative response of economic outcomes over time. In contrast, Highfill and Samuels [2026a], for example, takes an alternative measurement approach that classifies industries as AI intensive, and thus abstracts from relative intensities.

[3] Using the estimated real output semi elasticity of 0.1 to 0.2 percent for each 1 percentage point increase in the share of frequent AI users, a rise from 12 percent to 26 percent as documented in Makridis [2025] implies roughly 1.4 to 2.8 percent higher real output. Spread over the diffusion window from Q2:2024 to Q4:2025, this corresponds to about 0.9 to 1.8 percentage points of annualized real output growth. This should be interpreted as a temporary growth contribution during a period of rising adoption rather than a permanent increase in the long run growth rate.

2. Data and Measurement

Our analysis combines survey based measures of AI exposure with administrative and national accounts data on employment, output, and productivity. We link variation in AI use intensity across state industry cells and across industries to measures of economic performance over time.

2.1. AI Exposure Measures

We use two complementary measures of AI exposure. The first is based on quarterly data from the Gallup Workforce Panel and is constructed at the state by industry level spanning Q2:2025–Q1:2026. It captures the share of respondents in a given state-industry cell who report using AI frequently in their role. Gallup uses self-administered web surveys that ask over 20,000 respondents each quarter, "How often do you use artificial intelligence in your role," with response options ranging from daily to never. In our coding, "frequent" AI users are those who use AI daily or use a few times a week.[4] Before answering that question, respondents are given the following definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." Because this measure varies across both states and industries, it adds geographic variation in exposure and is well suited for our baseline design, which compares units within the same industry across different states.

The second measure is based on 2025 data from the Census Bureau's BTOS and varies only at the industry level [Bonney et al., 2024]. The BTOS is an ongoing, nationally- and industry-representative survey that is sent to over 200,000 employer businesses every two weeks and is published by North American Industry Classification System (NAICS) industry at the 2-digit level. In the 2025 questionnaire, BTOS asks firms: "In the last two weeks, did this business use Artificial Intelligence (AI) in any of its business functions?" and provides examples including machine learning, natural language processing, virtual agents, and voice recognition.[5] The Census defines AI as "computer systems and software that are able to perform tasks normally requiring human intelligence, such as decision-making, visual perception, speech recognition, and language processing." We compute an average quarterly AI usage share within each industry across fortnights to match the Gallup cadence. This yields an industry level exposure measure with no within-industry cross-state heterogeneity, so it supports a more aggregated empirical specification than the Gallup measure. These two measures allow us to compare results across a rich, state by industry exposure design based on worker reported frequent AI use and a coarse, industry-only exposure design based on higher frequency business reported AI use.

[4] https://www.gallup.com/699797/indicator-artificial-intelligence.aspx

[5] The question was phrased, "In the last two weeks, did this business use Artificial Intelligence (AI) in producing goods and services" when it was first asked in 2023 until late 2025.

2.2. LEHD Employment and Earnings Data

To measure labor market outcomes in the state by industry analysis, we use data from the Longitudinal Employer Household Dynamics (LEHD) program. The LEHD data provide annual measures of employment and earnings at the state by industry level. These data are especially useful for studying the labor market consequences of AI exposure because they offer a consistent panel of outcomes across detailed geographic and sectoral units. From the LEHD files, we construct log employment and log earnings. Employment captures the scale of labor demand within a state industry cell, while earnings capture changes in compensation. Because the LEHD data are available annually and align naturally with the state by industry structure of the Gallup AI exposure measure, they serve as the primary labor market outcomes in our baseline analysis.

A practical advantage of the LEHD data is that they permit the analysis of within industry cross state heterogeneity in labor market performance. This is central to our identification strategy, which relies on comparing state industry cells with different levels of AI exposure while controlling for time invariant differences across cells and common shocks to a given industry in a given year.

2.3. BEA State×Industry Data

We complement the LEHD labor market outcomes with BEA data on real output, providing annual inflation-adjusted gross domestic product (GDP) by state and industry, allowing us to examine whether greater AI exposure is associated not only with labor market changes but also with changes in production.[6] The BEA state by industry data are merged to the Gallup state by industry AI exposure measure and the LEHD outcomes using common state, industry, and year identifiers. We then construct log real output as an outcome variable in the baseline event study analysis. This measure is particularly valuable because it captures the productivity related dimension of AI adoption more directly than employment or earnings alone.

2.4. Sample Construction

Our estimation sample covers annual observations from 2016 through 2024. For the Gallup based analysis, the unit of observation is the state by industry by year cell for Q2:2025-Q1:2026.[7] For the BTOS based analysis, the unit of observation is the industry by year cell.

We restrict attention to cells that can be successfully matched across the relevant AI exposure, employment, earnings, wages, and output datasets. In the baseline Gallup analysis, observations are weighted by the number of survey respondents used to construct the AI exposure measure. This weighting gives greater influence to state industry cells for which the underlying survey based exposure measure is es-

[6] For the industry level BTOS analysis, we use several annual series from the BEA at the two digit NAICS level. These include employment, wages, value added, and real output.

[7] We aggregate across these quarters to create a sufficiently large sample at a state×industry level.

timated more precisely. More work must be done to obtain fully representative sampling, which is in progress, but we nonetheless believe these are a positive step forward [Makridis and Brynjolfsson, 2026].

All outcome variables are expressed in logarithms. This transformation facilitates interpretation in percentage terms and reduces the influence of extreme values.

2.5. Comparing the BTOS and Gallup

Figure 1 compares the evolution of AI use in the Gallup Workforce Panel and the Census Bureau's Business Trends and Outlook Survey [Bonney et al., 2024]. Both series show a clear upward trend, but the levels differ substantially. In the Gallup data, the share of workers reporting any AI use rises from roughly one fifth in mid 2023 to nearly one half by early 2026, while the share reporting frequent use also increases sharply, from about 10 percent to more than one quarter. The BTOS series also rises over the same period, but from a much lower base, increasing from roughly 4 percent to about 18 percent. These differences are informative rather than contradictory. Gallup measures worker reported use in a nationally representative sample of individuals, whereas BTOS measures firm reported use of AI in producing goods or services and other business functions.

The higher levels in Gallup are consistent with AI being used by workers for a broader set of tasks than those firms identify as core to production. The common upward trend across both series provides reassurance that they are capturing the same broad diffusion process, even if they operate at different levels of aggregation and reflect different margins of adoption. However, these results also suggest that BTOS may undercount the broader diffusion of AI, since firm reported use in production appears to capture a narrower margin of adoption than worker reported use in day to day tasks. One caveat is that the BTOS uses slightly different wording starting in Q4:2025. The original question asked whether the business used AI "in producing goods and services," while the revised 2025 questionnaire asks whether the business used AI "in any of its business functions." The concept is similar, but this wording change is worth noting when comparing periods. The share of "don't know" responses grows slightly from 7% to 11% and the BTOS share is still well below the total Gallup share for Q4:2025-Q1:2026.

One concern with the comparison thus far is that individuals may be using AI more frequently in their day-to-day tasks without firms having made deliberate organizational decisions to integrate AI into business processes. Worker-reported use rates could therefore overstate the degree to which AI has been structurally embedded in production workflows, even if they accurately capture individual behavior. To that end, we also draw on answers to the following question from the Gallup Workforce Panel: "To the best of your knowledge, has your organization begun integrating artificial intelligence (AI) technology or tools to improve organizational practices (e.g., to increase productivity, efficiency, and quality)?" This question captures the extensive margin of firm-level adoption as perceived by workers, providing an intermediate concept between individual task-level use and the firm-reported BTOS measure. Figure 2

plots the share of workers reporting organizational AI integration alongside the BTOS series. The organizational integration rate is still substantially higher than the BTOS measure throughout.

Figure 3 shows a strong positive cross industry relationship between the Gallup and Census BTOS measures of AI use (total). Industries with higher firm reported AI adoption in BTOS also tend to have higher worker reported AI use in the Gallup Workforce Panel, with a weighted correlation of 0.93. For nearly every industry, the Gallup measure is higher than the BTOS measure, which is consistent with the idea that worker reported AI use captures a broader set of task level applications than firm reported use in producing goods or services, as well as the possibility that worker level responses are more informative about aggregate firm behavior than a single person at the firm. Thus, the two datasets are measuring a common underlying diffusion process across industries, but the BTOS likely captures a narrower margin of adoption and may understate the breadth of AI use. Nevertheless, a disadvantage of the Gallup data is that, due to its relatively smaller sample size, may not be fully representative of all industries.

Table 1 summarizes the key conceptual differences between the two measures, which are substantial and bear directly on how the results should be interpreted. Whereas Gallup reports the share of individual workers using AI frequently, BTOS reports the share of employer businesses using AI in their business functions. These are not directly comparable even within the same industry, and no transformation of worker-level use rates into firm-level adoption rates is possible without within-firm deployment data that neither survey provides. The Gallup definition of AI—"a computer does or enhances a task normally performed by a person"—is intentionally broad and will capture casual, task-level use of general-purpose tools including large language models, whereas the BTOS question is oriented toward consequential, function-level deployment in business operations. While this breadth could explain some of the elevated Gallup worker use rates, the organizational integration measure—which asks workers specifically whether their firm has begun integrating AI into organizational practices rather than whether they personally use AI—remains substantially higher than BTOS even in these sectors, suggesting that the gap between worker-reported and firm-reported adoption cannot be attributed solely to casual individual use.

Finally, BTOS does not provide disaggregated public data at the state-by-industry level, so we cannot rule out that the absence of geographic granularity in the BTOS measure partly accounts for the weaker results it produces that will follow in the next section. We emphasize throughout that the Gallup-based specifications recover associations between industry-level worker adoption rates and economic performance, not firm-level adoption rates, and that the level gap between the two series reflects some conceptual differences in what is being measured rather than simple undercounting by either survey.

3. Empirical Strategy

To study the relationship between AI exposure and economic performance, we estimate event study specifications that relate differences in AI use intensity to changes in outcomes over time. Our goal is

to test whether sectors with higher AI exposure experienced differential changes in employment and and output in the years after the widespread diffusion of generative AI tools.

3.1. Gallup State×Industry AI Exposure

Our baseline specification uses a measure of AI exposure constructed at the state by industry level from Gallup survey data. Let y_{sit} denote an economic outcome for state s, industry i, and year t. The outcomes we study include log employment, log earnings, and log real output. We estimate:

$$y_{sit} = \alpha_{si} + \gamma_{it} + \sum_{\tau \neq 2020} \beta_\tau \left(\text{AIShare}_{si} \times \mathbf{1}\{t = \tau\}\right) + \varepsilon_{sit}, \tag{1}$$

where α_{si} are state by industry fixed effects, γ_{it} are industry by year fixed effects, and AIShare_{si} is the Gallup based measure of AI use intensity for state industry cell (s, i). The omitted year is 2020, so each coefficient β_τ is interpreted relative to 2020.

This specification compares changes over time across state industry cells with different levels of AI exposure, while absorbing all time invariant differences across state industry cells and all common shocks to a given industry in a given year. The inclusion of γ_{it} is especially important because it allows identification to come from within industry, across state differences in AI exposure rather than from aggregate industry trends alone. In the state by industry specifications, observations are weighted by the number of survey respondents used to construct the Gallup AI measure. Standard errors are clustered at the state by industry level.

The identifying assumption is that, absent differential AI exposure, state industry cells with higher and lower baseline AI use would have followed parallel trends in the outcomes of interest, conditional on the included fixed effects. The event study design provides a direct way to evaluate this assumption by examining pretrends in the years prior to the post 2020 period.

3.2. BTOS Industry Level AI Exposure

We also estimate an alternative specification using AI exposure derived from the BTOS data. In contrast to the Gallup measure, the BTOS measure varies only at the industry level. Let y_{it} denote an outcome for industry i in year t. We estimate:

$$y_{it} = \alpha_i + \delta_t + \sum_{\tau \neq 2020} \theta_\tau \left(\text{BTOSAI}_i \times \mathbf{1}\{t = \tau\}\right) + u_{it}, \tag{2}$$

where α_i are industry fixed effects, δ_t are year fixed effects, and BTOSAI_i is the industry level AI exposure measure. As before, 2020 is the omitted reference year.

This specification only contains cross industry variation in baseline AI exposure. Because the BTOS measure does not vary within industries across states, the model cannot include state by industry fixed

effects or industry by year fixed effects without absorbing the relevant variation. Identification therefore comes from tracing out differences in outcomes as a function of an industry's baseline AI exposure. Standard errors are clustered at the industry level. This second approach is very similar in nature to Highfill and Samuels [2026a].

3.3. Interpretation

The coefficients in both event study specifications should be interpreted as differential changes in outcomes for units with higher AI adoption relative to otherwise comparable units with lower AI adoption, normalized to 2020. In the specification using the Gallup Workforce Panel, a positive post-2020 coefficient means that state-industry cells with greater worker-reported AI use experienced faster growth in the outcome of interest after 2020, relative to the same industry in other states and relative to their own 2020 baseline. In the specification using the BTOS data, the interpretation is analogous, but the comparison is across industries rather than across states-industries.

An important feature of the design is that the AI use shares are measured in 2025 and Q1:2026 and then merged back to the panel of annual outcomes covering 2016–2024. The event study coefficients therefore trace out how states and industries that ranked higher in AI adoption as of the post period performed over time, answering the following question: are the states and industries with higher AI utilization also the ones exhibiting stronger economic performance across employment, output, and productivity? The design is well-suited to that question, but this timing structure makes it inherently difficult to interpret the estimates as causal. Because adoption is measured after the outcome period ends, reverse causality is a genuine concern: states and industries that experienced stronger output or productivity growth may have subsequently invested more in AI, generating a positive cross-sectional correlation between 2025 adoption rates and earlier economic performance that does not reflect AI's contribution to that performance. We cannot fully rule out this channel, and we note that the results should be read accordingly.

The main threat to identification beyond reverse causality is that AI adoption is not randomly assigned. High-adoption industries or states may differ systematically in ways that also predict faster output or employment growth, including differences in industrial composition, managerial quality, complementary intangible capital, or local demand conditions. The fixed effects structure is designed to address most of these concerns. State-by-industry fixed effects absorb all time-invariant differences across cells, such as persistent productivity gaps or industry location patterns, while industry-by-year fixed effects absorb shocks common to a given industry in a given year, including national technological changes, sector-specific demand shifts, and industry-wide business cycle movements. This makes the Gallup design especially demanding, since identification comes from within-industry differences across states rather than from broad sectoral trends alone.

The remaining identifying assumption is parallel trends: absent differential AI adoption, more and less AI-intensive units would have followed similar outcome paths after conditioning on the fixed effects. We assess the plausibility of this assumption by examining the pre-trends in the event study figures. The absence of large, growing pre-2020 coefficients, particularly for output, provides some support for this assumption, although the estimates should be read as evidence of differential associations rather than causal effects. By contrast, the BTOS specification is more vulnerable to omitted variable concerns because it does not contain state heterogeneity and thus cannot control for state×industry, industry×year, and state×year variation.

4. Main Results

Figure 4 shows a clear difference between employment and output responses in industries with higher Gallup based AI exposure. In Panel A, the coefficients for log employment are generally positive throughout the sample and become larger after 2020, peaking around 2022 before moderating somewhat in 2023 and 2024. The confidence intervals are wide as in Johnston and Makridis [2025], especially in the post period, so the employment results should be interpreted cautiously. Even so, the pattern does not support a simple displacement narrative in which greater AI exposure is associated with weaker labor demand. If anything, the point estimates suggest that higher exposure state industry cells experienced somewhat stronger employment expansions relative to 2020, although the uncertainty around those estimates remains substantial.[8]

Panel B shows a cleaner pattern for real output. The pre 2020 coefficients are slightly negative or near zero, while the post 2020 coefficients turn positive and remain positive through 2024. This shift is economically meaningful because the specification includes both state by industry fixed effects and industry by year fixed effects, so identification comes from within industry differences across states rather than from broad sectoral shocks. The two panels are thus consistent with a view of AI as an output-enhancing technology whose effects appear more clearly in output than in employment. The employment estimates do not suggest large immediate labor market losses, while the output estimates indicate that more AI exposed state industry cells saw stronger growth after 2020. That combination is consistent with AI being associated with efficiency gains and possibly expanded production, without an equally sharp short-run contraction in labor demand, though the mechanisms driving this pattern cannot be established from the current design.

To address concerns that the effects on real output are not informative for productivity, Figure 5 reports the results associated with a similar event study, but this time using a simplified measure of TFP as an outcome variable. In particular, we residualize log real output with log employment (also state ×

[8] As a robustness check, we also used 2019 as the reference year and the results were typically either the same or showed stronger effects.

industry) and log private fixed assets (only industry). We find similar results.[9],[10] In contrast, Figure 6 reports the results when using BTOS at an industry level. There is slight evidence of an upward trajectory, but not enough variation exists to estimate meaningful effects.

Recognizing that the small sample size in some industry-state cells from the Workforce Panel could lead to spurious correlations, Figure 7 replicates the baseline specification after restricting to cells with at least 30 respondents determining the share of frequent AI users. The results largely hold for employment and real output, but the patterns for TFP become null.

These estimates imply effects that are economically consequential once mapped into the observed pace of diffusion in AI use. Using our real output semi elasticity, the increase in the share of frequent AI users from roughly 12 percent in Q2:2024 to 26 percent in Q4:2025 would correspond to about 1.4 to 2.8 percent higher real output. Spread over that diffusion window, this back of the envelope calculation implies roughly 0.9 to 1.8 percentage points of annualized real output growth, although this should be interpreted as a growth contribution during a period of rising adoption rather than a permanent increase in the long run growth rate. Relative to recent macroeconomic benchmarks, this places our reduced form estimates above the more conservative aggregate calibration in Acemoglu [2024] and closer to the more expansionary range discussed by Aghion and Bunel [2024], while still not being directly comparable because our estimates are differential state by industry associations rather than an aggregate TFP projection.

Although the event study estimates show no evidence against parallel trends, we introduce an additional robustness exercise that instruments realized AI use with predetermined exposure to generative AI. Following Johnston and Makridis [2025], we construct a state-industry GenAI exposure index using occupation level exposure scores from Eloundou et al. [2024] aggregated with American Community Survey occupation shares from 2015 to 2019. Because these shares are measured before the diffusion of enterprise generative AI tools, the instrument is anchored in pre existing occupational composition rather than post treatment adjustments in employment or task allocation. We then interact this pre pandemic exposure index with year indicators and use those interactions as instruments for the corresponding interactions between year indicators and realized AI use in the baseline event study specification. Intuitively, this design asks whether state-industry cells with greater ex ante exposure, based on their pre pandemic task mix, experienced larger post 2020 changes in outcomes through their induced increases in AI use. Given state-industry and industry-year fixed effects, identification comes from differential changes within industries across states and within states across industries, rather than from sectoral or

[9] We also examined whether growth in AI use at a state×industry level predicts higher economic activity. While the Gallup Workforce Panel is not yet large enough to create within state-industry growth rates, and thus many cells are too small to include in the analysis, we do find statistically significant effects on log employment following the introduction of ChatGPT for states and industries with higher growth in AI use.

[10] We also distinguish between producers of AI, compared with users. Data on the proportion of AI-classified jobs, relative to the total, is less than 2%; see Andreadis et al. [2025] for evidence on the determinants of AI job shares across counties.

macroeconomic shocks.

Figure 8 shows that IV event study results are broadly consistent with the main findings for output, earnings, and productivity, while the employment results differ and exhibit large confidence intervals. Panel A shows that the coefficients for log employment are noisy throughout the sample and do not display the same clear post 2020 pattern as in the baseline estimates. The point estimates are generally positive in several years, but the confidence intervals are wide and include economically modest negative effects as well, which suggests that the employment response is not robust to using pre period generative AI exposure as an instrument for realized AI use.

By contrast, the remaining panels are more supportive of the baseline interpretation. Panel B indicates that log earnings, which were somewhat less stable in the baseline specifications, now display a clearer positive post 2020 pattern, especially from 2021 onward. Although precision remains limited in some years, the earnings results are now directionally consistent with the broader view that AI adoption is associated with improvements in labor market outcomes rather than wage compression. Panel C shows the strongest evidence for real output: coefficients rise after 2020 and remain positive through the end of the sample, with magnitudes that are consistent with meaningful gains in production among state industry cells with greater instrumented AI utilization. Panel D shows a similar pattern for total factor productivity, where the post 2020 coefficients become positive and remain elevated, suggesting that the output gains are not driven solely by input accumulation but also reflect improvements in efficiency. In this sense, the IV results are consistent with the interpretation that associations between AI use and economic performance operate primarily through productivity and output, possibly with associated gains in earnings.

5. Intangible Capital and Long-run Implications

Our results provide suggestive evidence that AI adoption is linked with improvements in productivity already, but there are a handful of reasons to suspect that the productivity benefits will be even larger in the longer-run because AI is a general purpose technology whose full value is unlocked only when complementary investments have been made [Brynjolfsson et al., 2021] and production structures have been reorganized to make efficient use of AI. To explore the relationship between AI, complementary investment, and production structure, we use BEA data to analyze how AI interacts with R&D production and how AI intensity relates to the use of intangible capital. While previous examined the relationship between AI-intensity and the cost structure of industries [Highfill and Samuels, 2026b], it did not focus specifically on the intangible economy.

Our focus on the relationship between AI and intangible capital is motivated by AI as a potential driver of structural change in the information and knowledge economy. World Intellectual Property Organization and Luiss University [2025] describes the ongoing shift to intangible investment and the particularly

important role of intangible capital in advanced economies. Yet, the relationship between AI use and intangible capital is not clearly predicted by economic theory and may evolve over time. For instance, if an AI platform is a substitute for traditional software, AI use may be associated with a decline in software use. Importantly, because AI use is not directly measured as an input, models of production that exclude AI as an input, could incorrectly conclude that software use was declining in importance when in reality AI was serving as an unmeasured substitute. Alternatively, general purpose AI, which is functionally just software, may lead to additional new purchases of software investment as AI gets embedded in many software products. Similarly, AI may reduce the costs of in house R&D, thus lead to an observed decline in investment in *measured* R&D even if the real output of R&D is unchanged, because the nominal value of R&D is measured at cost. Conversely, AI could provide a boon to R&D investment by making each researcher more productive and raising the rate of return on R&D investment.

Figure 9 examines whether higher AI use is associated with shifts in the composition of intangible capital across industries. The results are largely null. Arts and Entertainment Originals IP capital shares (Panel A) show no systematic pre- or post-2020 pattern, with coefficients clustered near zero throughout the sample. R&D shares (Panel B) display negative pre-2020 coefficients that recover toward zero by 2021 before rising modestly in 2022 and 2023, but the confidence intervals are wide enough that no post-2020 trend can be distinguished from noise. Software shares (Panel C) follow a similar trajectory—positive early in the sample, declining through 2021–2022, and turning negative by 2023—again without statistical precision sufficient to support a directional claim. Data capital shares (Panel D) are uniformly negative across the full sample, with the largest negative point estimates in 2016–2017, before converging toward zero by 2022–2023.[11] The four panels provide no robust evidence that industries with higher AI utilization have systematically reallocated their intangible capital portfolios. The R&D panel shows the most directionally consistent post-2020 pattern, which aligns with theoretical predictions, but the signal is weak and the mixed results across panels argue for caution. However, AI use at the industry level may not yet register in official capital account data—consistent with the broader measurement challenges documented in Highfill and Samuels [2026a] and Makridis and Brynjolfsson [2026]. An additional concern is reverse causality: industries investing heavily in R&D and intangible capital prior to 2020 may have been better positioned to adopt AI, which would produce correlations between the 2025 Gallup measure and pre-2020 capital patterns without implying that AI drove those patterns. The intangibles analysis is therefore better read as exploratory, consistent with the J-curve hypothesis of Brynjolfsson et al. [2021], rather than as evidence of AI-driven capital reallocation. The null results could also reflect the absence of the within-industry geographic variation that sharpens identification in the baseline employment and output results.

We now present these results in light of the measurement framework developed by Corrado et al. [2005] and Corrado et al. [2009]. The official national accounts which we draw upon include a large portion of

[11] The inclusion of data capital is consistent with the changes recommended by the 2025 system of national accounts, but not yet included in the official data. We have taken our measures of data as an asset from Samuels et al. [2026].

the intangible assets discussed in those papers, but exclude some other aspects, such as organizational capital. The four capital share outcomes we examine map into the CHS taxonomy, yet none displays a robust post-2020 response to higher AI utilization. This is consistent with a central implication of Corrado et al. [2021], who model AI as a combination of measured intangibles (software) and unmeasured intangibles (databases and organizational routines), and find that AI-specific investment is not separately identifiable with precision in standard national accounts because some of it is subsumed within broader software and R&D aggregates. If the productivity-enhancing components of AI investment are concentrated in the unmeasured portion of the intangible capital stock, then regressions using capital share data will recover attenuated or null estimates even when real reallocation is occurring. The absence of a detectable signal in our industry-level specifications is therefore not strong evidence against AI-driven capital deepening; it more plausibly reflects the asset boundary of official statistics.

Corrado et al. [2024] specifically focus on data capital, showing that it contributes roughly half of measured intangible capital across nine European economies. That mechanism offers a specific channel through which AI adoption could generate the output and productivity gains documented in Figure 4 and Figure 5—operating through data-intensive production processes—without producing a detectable shift in official capital shares. To address this, in the results given in Figure 9, we have included data capital services based on the work in Samuels et al. [2026]. Based on those working estimates, AI intensity is not associated with a shift toward data assets.

Corrado et al. [2022] further document that intangible investment as a share of GDP has plateaued in many advanced economies since the mid-2010s, which implies that post-2020 AI-driven investment may not yet be large enough in aggregate to shift industry-level capital shares within the 2016–2023 window we study. Detecting AI's footprint in capital accounts will therefore require either expanding the measures to include AI-specific investment categories, as Makridis and Brynjolfsson [2026] propose, or combining administrative capital data with the granular, worker-reported utilization measures. We leave these aims for future research.

6. Conclusion

This paper studies whether AI utilization is associated with changes in output and labor market outcomes across U.S. industries. Using a new state by industry measure of worker reported AI use from the Gallup Workforce Panel and an industry level measure of firm reported AI use from the Census Bureau's Business Trends and Outlook Survey, we link realized adoption to administrative and national accounts data from LEHD and BEA over 2016 to 2024. Our contribution is largely empirical: rather than relying only on occupational exposure proxies or broad technological classifications, we focus on actual AI use. That distinction matters. Exposure measures are informative about where AI could matter, but realized use is closer to the margin through which production processes, organizational routines, and worker engagement actually change. Across specifications, the associations between AI utilization

and real output and productivity are more consistently positive than those between AI utilization and employment, with the worker based Gallup measure generally recovering sharper patterns than the more aggregated BTOS measure.

The findings reinforce the case for a new measurement agenda around AI. One of the main lessons is that the economic effects of AI are easier to detect when researchers can move beyond coarse proxies and observe actual use at a sufficiently granular level. A natural next step is therefore to combine telemetry based measures of software interaction, platform activity, and workflow integration with carefully designed respondent level data on how, when, and for what tasks workers and firms use AI. For example, instead of relying on surveys of AI use, utilization can be observed directly using tokens available from AI software providers that track the volume of model outputs generated when AI is used [New et al., 2025]. Makridis and Brynjolfsson [2026] argue that progress in economic measurement may depend on integrating new forms of AI investment and use data into official statistics, rather than forcing AI into categories built for earlier vintages of capital. That combination of behavioral data and survey based context may be essential for distinguishing experimentation from sustained adoption, measuring complementary intangible investments, and tracing how AI changes the organization of production over time.

Nevertheless, several limitations point to directions for future work. First, the available measures of AI use from the Gallup Workforce Panel are limited by sample size, which stifles the ability to measure within cell changes in adoption over time and does not represent all industries. However, the sample continues to grow, especially now that it is quarterly. Second, even with rich fixed effects, identification remains challenging because AI adoption is endogenous to managerial quality, organizational capital, and local demand conditions that may not be fully observed, although we have tried to address that limitation using generative AI exposure prior to the pandemic. To that end, more work is needed to identify potential instruments that have a strong first-stage and yet remain plausibly exogenous. Third, the results are most informative about average relationships at the state industry level, not about the underlying mechanisms within firms or across occupations. Future work should therefore connect survey based measures of worker and firm use to administrative records, platform telemetry, and firm level production data to better isolate causal effects, distinguish augmentation from automation, and understand how the productivity gains from AI are distributed across workers, firms, and places. In that sense, this paper should be read as an early step in a broader measurement project: the evidence is consistent with AI being associated with stronger output and productivity growth at the state-industry level, though identifying the causal contribution of AI specifically remains difficult given the endogeneity of adoption and the nascent state of AI diffusion before 2022.

References

Acemoglu, D. (2024). The simple macroeconomics of ai. Working Paper 32487, National Bureau of Economic Research.

Aghion, P. and Bunel, S. (2024). Ai and growth: Where do we stand?

Andreadis, L., Kalotychou, E., Chatzikonstantinou, M., Louca, C., and Makridis, C. A. (2025). Local heterogeneity in artificial intelligence jobs over time and space. *American Economic Association: Papers and Proceedings*, 115:29–34.

Bick, A., Blandin, A., and Deming, D. J. (2025). The rapid adoption of generative ai. NBER Working Paper 32966, National Bureau of Economic Research.

Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., and Savage, K. (2024). Tracking firm use of ai in real time: A snapshot from the business trends and outlook survey. NBER Working Paper 32319, National Bureau of Economic Research.

Brynjolfsson, E., Chandar, B., and Chen, R. (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Corrado, C., Haskel, J., Iommi, M., Jona-Lasinio, C., and Bontadini, F. (2024). Data, intangible capital, and productivity. In *Technology, Productivity, and Economic Growth*. University of Chicago Press, Chicago.

Corrado, C., Haskel, J., and Jona-Lasinio, C. (2021). Artificial intelligence and productivity: An intangible assets approach. *Oxford Review of Economic Policy*, 37(3):435–458.

Corrado, C., Haskel, J., Jona-Lasinio, C., and Iommi, M. (2022). Intangible capital and modern economies. *Journal of Economic Perspectives*, 36(3):3–28.

Corrado, C., Hulten, C. R., and Sichel, D. E. (2005). Measuring capital and technology: An expanded framework. In Corrado, C., Haltiwanger, J., and Sichel, D., editors, *Measuring Capital in the New Economy*, volume 65 of *Studies in Income and Wealth*, pages 11–46. University of Chicago Press, Chicago.

Corrado, C., Hulten, C. R., and Sichel, D. E. (2009). Intangible capital and U.S. economic growth. *Review of Income and Wealth*, 55(3):661–685.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: An early look at the labor market impact potential of large language models. *Science*, 384(6702):1306–1308.

Highfill, T. and Samuels, J. D. (2026a). Early estimates of the impact of ai within bea's industry economic accounts. BEA Working Paper WP2026-3, U.S. Bureau of Economic Analysis. BEA Working Paper Series.

Highfill, T. and Samuels, J. D. (2026b). Early evidence on the relationship between ai, costs, and prices within bea's industry economic accounts. Bea working paper, U.S. Bureau of Economic Analysis. Forthcoming, BEA Working Paper Series.

Highfill, T., Wasshausen, D., and Prunchak, G. (2025). Concepts and challenges of measuring production of artificial intelligence in the u.s. economy. BEA Working Paper WP2025-1, U.S. Bureau of Economic Analysis. BEA Working Paper Series.

Humlum, A. and Vestergaard, E. (2025). Large language models, small labor market effects. BFI Working Paper 2025-56, Becker Friedman Institute for Economics at the University of Chicago.

Johnston, A. and Makridis, C. A. (2025). Ai, output, and employment. *SSRN working paper*.

Makridis, C. A. (2025). Organizational transmission of ai: Managerial influence on generative ai adoption. *SSRN working paper*.

Makridis, C. A. and Brynjolfsson, E. (2026). Counting ai: A blueprint to integrate ai investment and use data into us national statistics. Technical report, Brookings Institution.

New, J., Meyjes, M., and Carson, A. (2025). The u.s. needs a generative ai intensity index. Technical report, SeedAI.

Samuels, J. D., Santiago Calderón, J. B., and Garner, C. (2026). The impact of capitalizing data on productivity growth in the u.s. BEA Working Paper WP2026-7, U.S. Bureau of Economic Analysis. BEA Working Paper Series.

World Intellectual Property Organization and Luiss University (2025). World intangible investment highlights 2025: Better data for better policy. Technical report, World Intellectual Property Organization (WIPO), Geneva, Switzerland. Co-published with Luiss Business School.

Tables and Figures

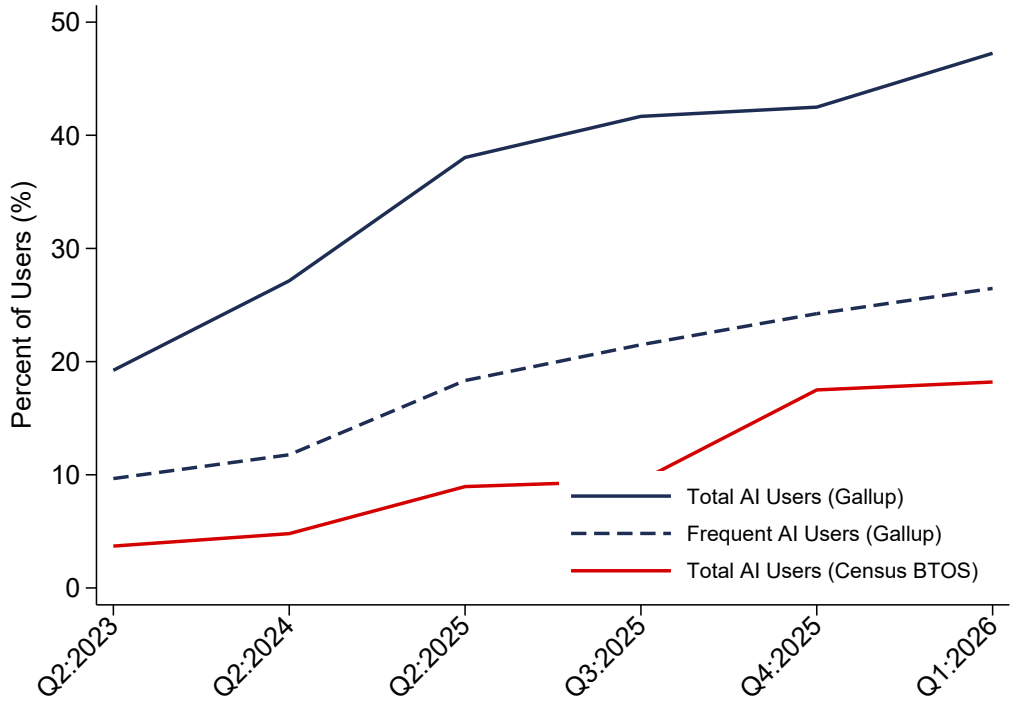


Figure 1. Comparing the Gallup and Census AI Measures Over Time

Notes.—Sources: Gallup Workforce Panel and Census Bureau Business Trends and Outlook Survey (BTOS). The figure compares alternative measures of AI use over time. The solid navy line reports the share of Gallup Workforce Panel respondents who report any AI use in their role, defined as the sum of respondents reporting frequent use (daily or multiple times a week) and occasional use (sometimes during a year). The dashed navy line reports the share of respondents who report frequent AI use. The red line reports the share of firms in the Census Bureau's Business Trends and Outlook Survey that report using AI in producing goods or services. The Gallup series is based on a nationally representative set of workers (using survey weights), while the BTOS series is based on a survey of employer businesses and therefore reflects firm reported adoption rather than worker reported use. The BTOS measure is shown for the aggregate sector only. To align the two series in a common visual format, the BTOS observation labeled Q2:2023 corresponds to the Q3:2023 BTOS release and is shifted back one quarter in the figure.

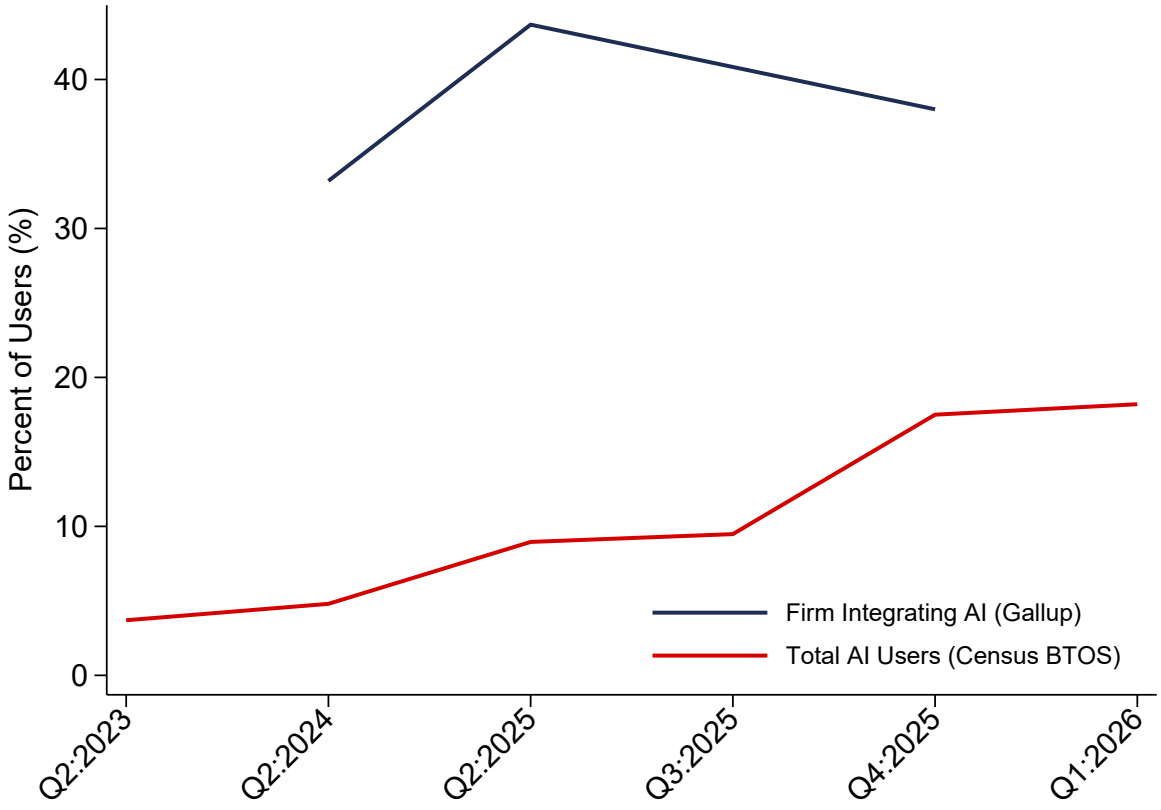


Figure 2. Worker Perceptions versus Firm-Reported Adoption

Notes.—Sources: Gallup Workforce Panel and Census Bureau Business Trends and Outlook Survey (BTOS). The figure compares two measures of AI adoption over time. The solid navy line reports the share of Gallup Workforce Panel respondents who answer yes to the following question: "To the best of your knowledge, has your organization begun integrating artificial intelligence (AI) technology or tools to improve organizational practices (e.g., to increase productivity, efficiency, and quality)?" This measure captures the extensive margin of firm-level AI adoption as perceived by workers. The solid red line reports the share of firms in the Census Bureau's Business Trends and Outlook Survey that report using AI in any of their business functions (revised 2025 wording; previously asked whether the business used AI in producing goods or services). The Gallup series is based on a nationally representative sample of workers using survey weights, while the BTOS series is based on a survey of employer businesses and reflects firm-reported adoption rather than worker-reported perceptions. The BTOS measure is shown for the aggregate sector only. To align the two series in a common visual format, the BTOS observation labeled Q2:2023 corresponds to the Q3:2023 BTOS release and is shifted back one quarter in the figure.

Table 1. Comparing AI Measures: Gallup Workforce Panel and Census BTOS

Dimension	Gallup: Worker Use	Gallup: Org. Integration	Census BTOS
Respondent	Individual worker	Individual worker	Firm representative
Unit of measure	Share of workers using AI frequently	Share of workers reporting their organization has begun integrating AI	Share of firms using AI in business functions
Survey question	"How often do you use artificial intelligence in your role?" Frequency options (daily, few times a week, sometimes, never) allow distinction between intensive and extensive margins of use.	"Has your organization begun integrating AI technology or tools to improve organizational practices?" Captures extensive margin only — binary yes/no on organizational adoption.	"Did this business use AI in any of its business functions?" (revised 2025); previously "in producing goods or services." Captures extensive margin only — binary yes/no on firm-level use.
Definition of AI	"A computer does or enhances a task normally performed by a person" — broad, encompasses casual and incidental use	Same as worker use	"Computer systems able to perform tasks normally requiring human intelligence"; examples include ML, NLP, virtual agents, voice recognition
Margin captured	Individual task-level, including casual use of general-purpose tools	Organizational-level adoption decision as perceived by worker	Function-level deployment in business operations
Survey frequency	Quarterly (Q2:2025 onward); annual prior	Quarterly (Q2:2025 onward)	Fortnightly, aggregated to quarterly
Sample size	∼20,000 workers per quarter, nationally representative	Same	>200,000 employer businesses
Comparability to firm adoption	Not directly comparable — requires within-firm deployment information unavailable in either survey	Closer to firm-level concept but filtered through worker perception	Direct firm-level measure

Notes.—Sources: Gallup Workforce Panel and Census Bureau Business Trends and Outlook Survey (BTOS). The three measures capture different margins of AI adoption and are not directly comparable in levels. The Gallup worker use measure is the primary variable in the baseline event study specifications, using the frequent use category (daily or a few times a week) to capture the intensive margin of adoption. The organizational integration measure provides an intermediate concept between individual task-level use and firm-level deployment. The BTOS measure varies only at the 2-digit NAICS industry level and does not permit state-by-industry identification. The high cross-industry correlation between Gallup worker use and BTOS (0.93, Figure 3) indicates the two series track the same broad diffusion process, but the level gap of 15–25 percentage points reflects genuine conceptual differences rather than simple undercounting by either survey.

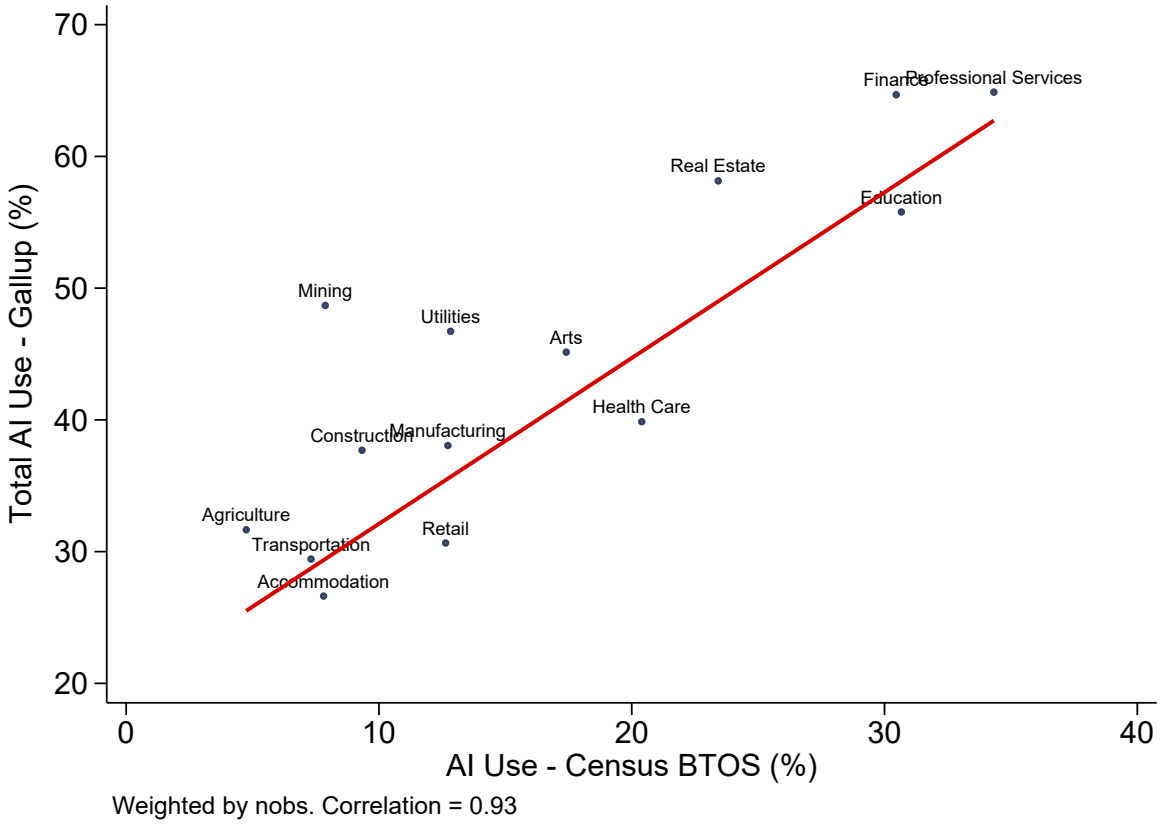
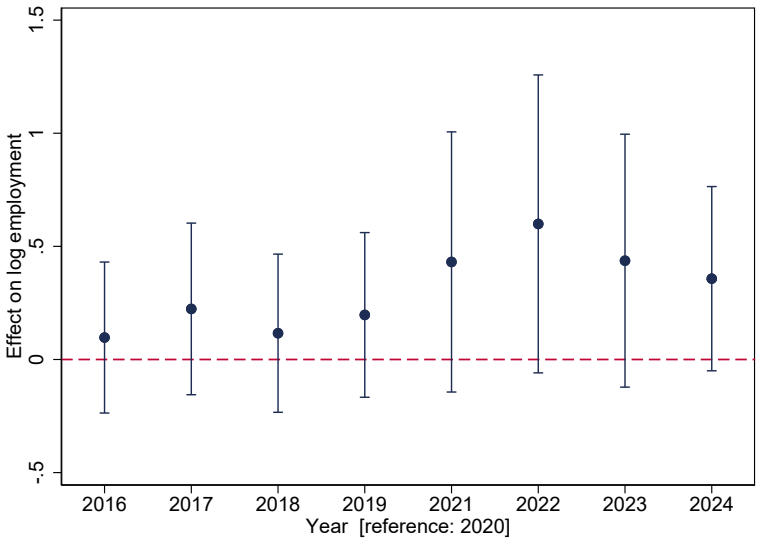


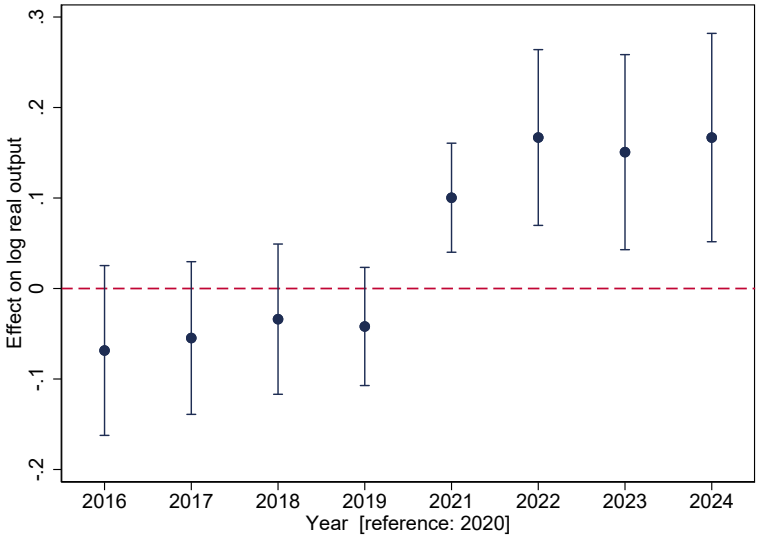
Figure 3. Comparing the Gallup and Census AI Measures Across Industries

Notes.—Sources: Gallup Workforce Panel and Census Bureau Business Trends and Outlook Survey (BTOS). The figure reports the cross-industry relationship between the Gallup and BTOS measures of AI use using the most recent matched data (Q4:2025 and Q1:2026 averaged together after the BTOS updated its wording). Each point corresponds to a two-digit NAICS industry, with labels indicating the industry name. The horizontal axis reports the BTOS share of firms in an industry that report using AI in producing goods or services, while the vertical axis reports the corresponding Gallup share for the same industry. The fitted red line is estimated using weights based on the Gallup cell size, and the note reports the weighted correlation between the two measures. The Gallup series is based on a nationally representative sample of workers and captures worker reported AI use, while the BTOS series is based on a survey of employer businesses and captures firm reported use in production.

Figure 4. Baseline Event Study Results Using Gallup AI Shares



a. Employment



b. Real Output

Notes.—Source: Gallup survey data, Longitudinal Employer-Household Dynamics (LEHD), and Bureau of Economic Analysis (BEA), 2016–2024. The figure reports the coefficients associated with regressions of log employment and log real output on the Gallup based AI share measure interacted with year fixed effects, conditional on state by industry and industry by year fixed effects. Panel A uses LEHD employment data and Panel B uses BEA state industry real output data. Coefficients are normalized relative to 2020, the omitted reference year. Observations are weighted by the number of Gallup respondents used to construct the AI share measure. Standard errors are clustered at the state by industry level.

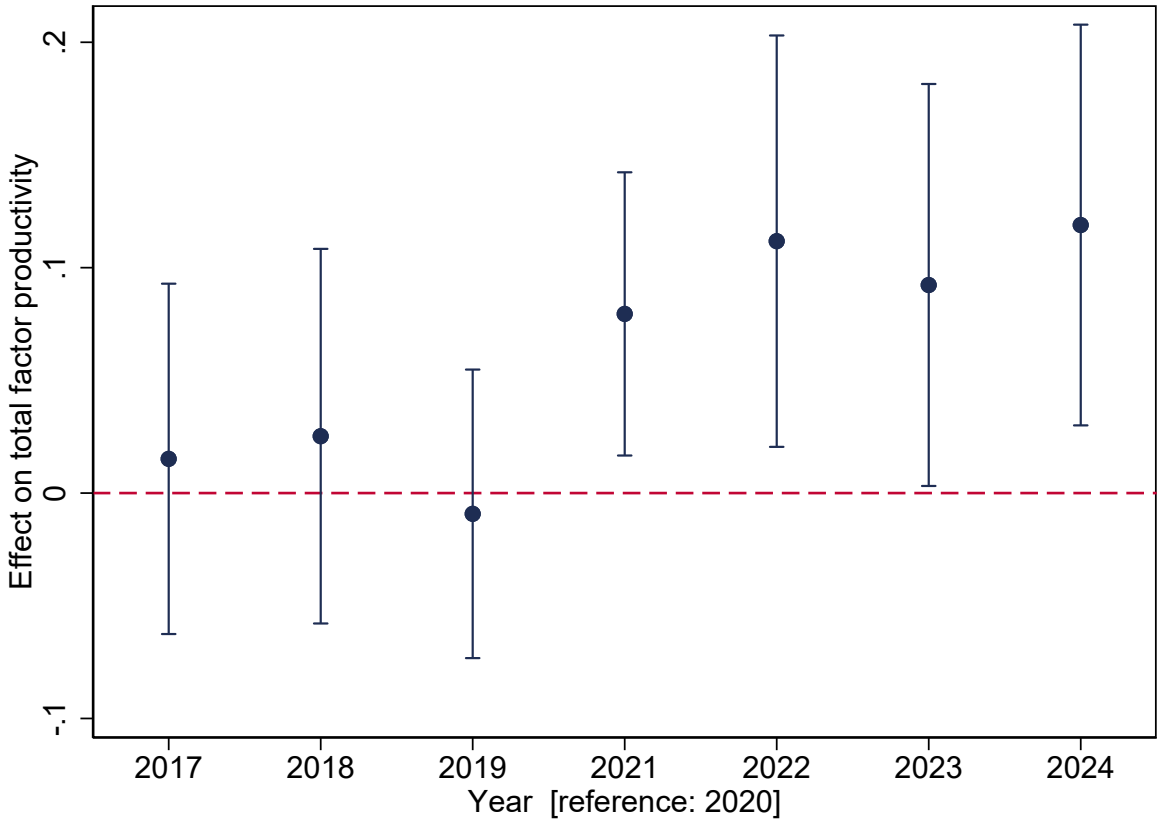
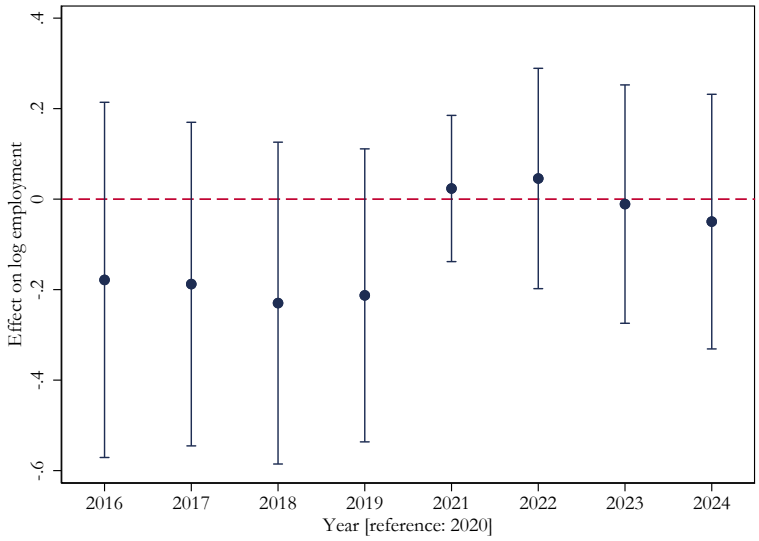


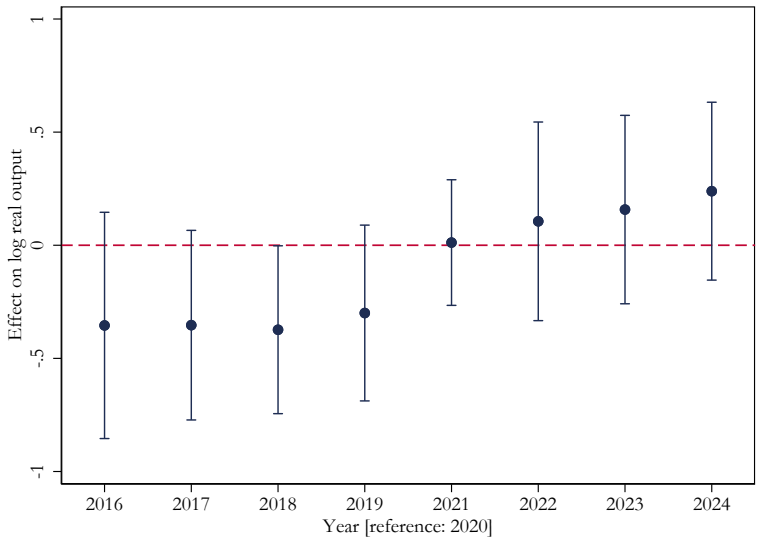
Figure 5. Event Study Robustness on Total Factor Productivity

Notes.—Sources: Gallup survey data, Longitudinal Employer-Household Dynamics (LEHD), and Bureau of Economic Analysis (BEA), 2016–2024. The figure reports the coefficients associated with regressions of total factor productivity (TFP) on the Gallup based AI share measure interacted with year fixed effects, conditional on state by industry and industry by year fixed effects. TFP is measured by taking the residual from a regression of log real output on log employment and log private assets (only at the NAICS2 level). Coefficients are normalized relative to 2020, the omitted reference year. Observations are weighted by the number of Gallup respondents used to construct the AI share measure. Standard errors are clustered at the state by industry level.

Figure 6. Baseline Event Study Results Using BTOS AI Shares



a. Employment

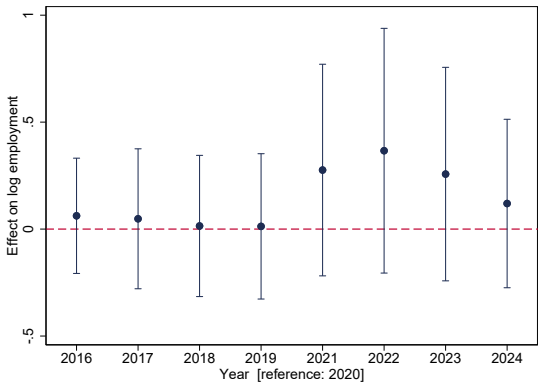


b. Real Output

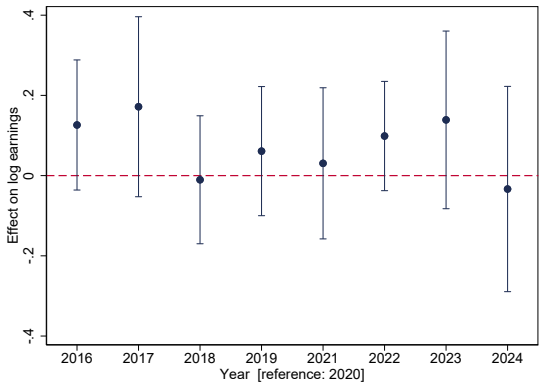
Notes.—Source: Census Business Trends and Outlook Survey (BTOS) and Bureau of Economic Analysis (BEA), 2016–2024. The figure reports the coefficients associated with regressions of log employment and log real output on the BTOS based AI share measure interacted with year fixed effects, conditional on NAICS2 and year fixed effects. Panel A uses BEA industry employment data and Panel B uses BEA industry real output data. The BTOS AI share measure is constructed at the NAICS2 industry level and does not include state level heterogeneity. Coefficients are normalized relative to 2020, the omitted reference year. Standard errors are clustered at the NAICS2 industry level.

Figure 7. Robustness Restricting Baseline to Larger State×Industry Cells

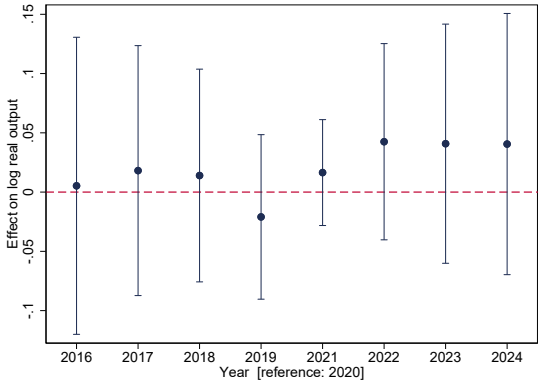
a. Panel A: log(Employment)



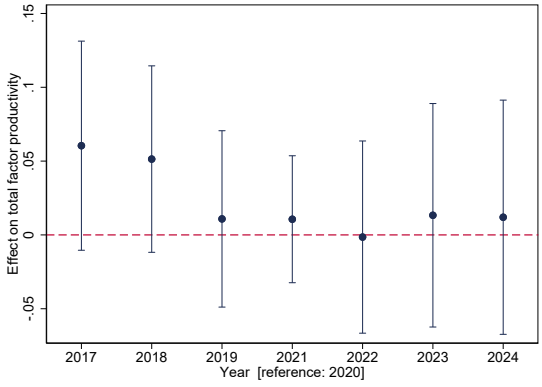
b. Panel B: log(Earnings)



c. Panel C: log(Real Output)



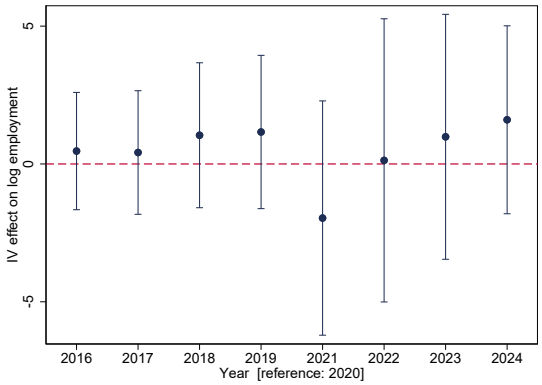
d. Panel D: Total Factor Productivity



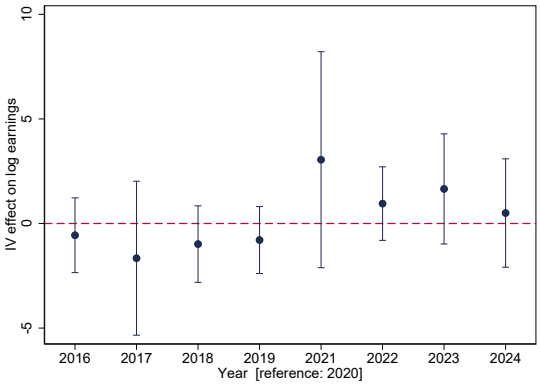
Notes.—Source: Gallup survey data, Longitudinal Employer-Household Dynamics (LEHD), and Bureau of Economic Analysis (BEA), 2016–2024. The figure reports the coefficients associated with regressions of log employment, log earnings, log real output, and total factor productivity (TFP) on the Gallup based AI share measure interacted with year fixed effects, conditional on state by industry and industry by year fixed effects, and restricting the sample to industry-state cells with at least 30 respondents from the Workforce Panel. TFP is measured by taking the residual from a regression of log real output on log employment and log private assets (only at the NAICS2 level). Coefficients are normalized relative to 2020, the omitted reference year. Observations are weighted by the number of Gallup respondents used to construct the AI share measure. Standard errors are clustered at the state by industry level.

Figure 8. Robustness Using Generative AI Exposure to Instrument for AI Utilization

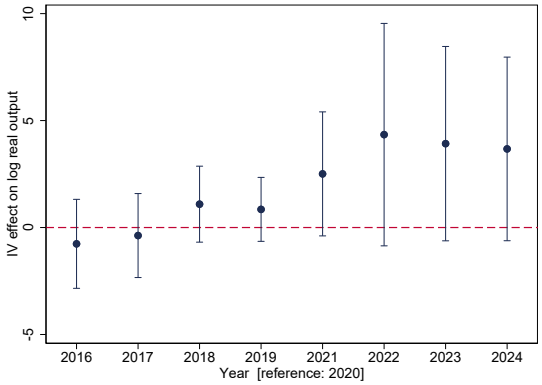
a. Panel A: log(Employment)



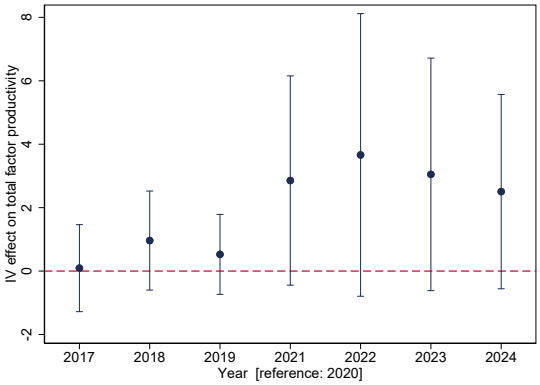
b. Panel B: log(Earnings)



c. Panel C: log(Real Output)



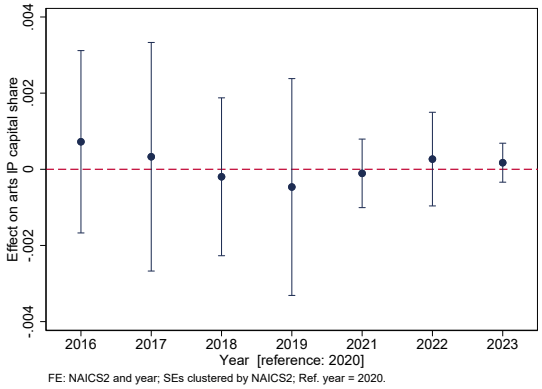
d. Panel D: Total Factor Productivity



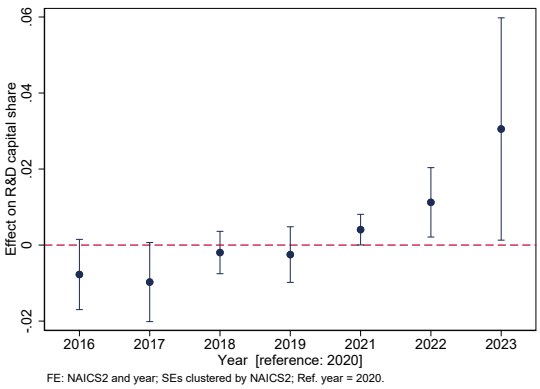
Sources: Gallup survey data, Longitudinal Employer-Household Dynamics (LEHD), Bureau of Economic Analysis (BEA), Eloundou et al. [2024], and the American Community Survey (ACS), 2016–2024, with ACS data from 2015–2019 used to construct the instrument. The figure reports instrumental variables event study coefficients from regressions of log employment, log earnings, log real output, and total factor productivity (TFP) on the Gallup based AI use share interacted with year indicators, where each year specific interaction is instrumented with the corresponding interaction between year indicators and a pre period state-industry exposure measure to generative AI. Following Johnston and Makridis [2025], the exposure measure is constructed by aggregating the occupation level LLM exposure scores in Eloundou et al. [2024] using ACS occupation shares within state-industry cells over 2015–2019. All specifications include state by industry and industry by year fixed effects. Coefficients are normalized relative to 2020, the omitted reference year. Observations are weighted by the number of Gallup respondents used to construct the AI use share, and standard errors are clustered at the state by industry level. TFP is measured as the residual from a regression of log real output on log employment and log private assets.

Figure 9. Dynamic Effects of AI Use on Intangible Capital Shares

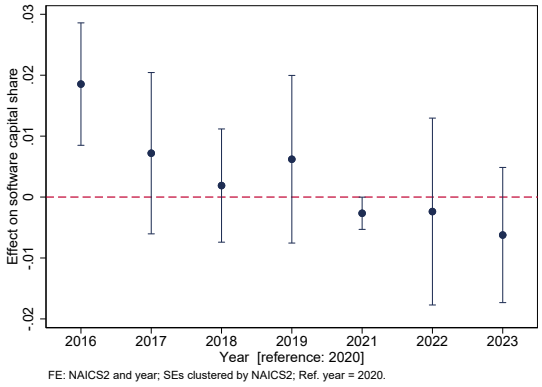
a. Panel A: Arts IP Capital Share



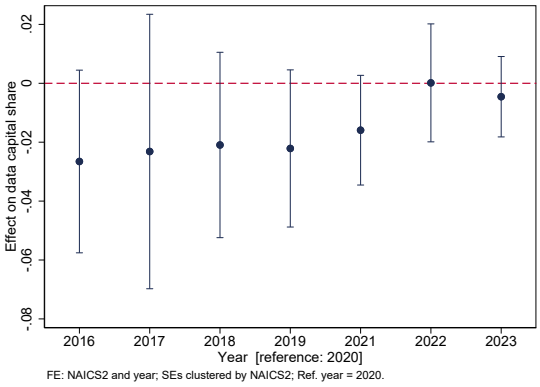
b. Panel B: R&D Capital Share



c. Panel C: Software Capital Share



d. Panel D: Data Capital Share



Notes.—Sources: Gallup Workforce Panel (2025); Integrated Labor Productivity Accounts (ILPA), 2016–2023. The figure reports coefficients from regressions of the industry share of arts IP, R&D, software, and data capital on the Gallup-based measure of frequent AI use, interacted with year indicators and normalized to 2020, conditional on NAICS2 and year fixed effects. The AI use measure is collapsed to the NAICS2 level using data from 2025 and merged to annual ILPA data. Observations are weighted by the number of Gallup survey respondents used to construct the AI use measure. Standard errors are clustered at the NAICS2 level.