BEA Working Paper Series, WP2026–18

AI Utilization and Economic Performance

Author Christos A. Makridis, Gallup, Arizona State University, Stanford University;
Tina Highfill and Jon D. Samuels, U.S. Bureau of Economic Analysis[1]

Contact cmakridi@asu.edu

Date August 2026

Abstract Do state-industry cells that report more intensive artificial intelligence (AI) use in 2025–2026 also exhibit stronger economic performance over 2016–2024? We combine a state-by-industry measure of worker-reported frequent AI use from the Gallup Workforce Panel with an industry-level measure of employer-reported AI use from the U.S. Census Bureau's Business Trends and Outlook Survey (BTOS), pooling observations from Q2:2025 through Q1:2026. We link these measures to annual employment, earnings, and real-output data from the Longitudinal Employer-Household Dynamics program and the U.S. Bureau of Economic Analysis and estimate retrospective dynamic specifications with state×industry and industry×year fixed effects. State-industry cells with higher worker-reported AI use exhibit stronger post-2020 real-output paths and positive, though imprecisely estimated, employment differences. The industry-only BTOS specification does not recover the same output pattern. These estimates are descriptive rather than causal and underscore the value of granular measures of realized AI use.

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

JEL codes O33, J23, J24, D22, E24, C23

[1] Thank you to Erik Brynjolfsson, Emin Dinlersoz, Jim Harter, Andy Kemp, and Sabrina Yu for comments.

Suggested citation:
Makridis, Christos A., Tina Highfill, and Jon D. Samuels. *AI Utilization and Economic Performance.* Working paper no. WP2026–18. Washington, DC: U.S. Bureau of Economic Analysis, August 2026, https://doi.org/10.66137/HMXC8574.

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. [2026a], Makridis [2026]. Yet the magnitude and timing of its macroeconomic costs and benefits remain uncertain and widely debated [Aghion and Bunel, 2024, Acemoglu, 2025, Bick et al., 2026b, Filippucci et al., 2026, Johnston and Makridis, 2026, Yotzov et al., 2026]. Public attention to AI accelerated sharply after 2022, but the relevant economic adjustments began earlier.[2] Unlike traditional capital investments, such as machinery or information technology (IT) 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 a technical challenge: it shapes how quickly economists and policymakers 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 U.S. Census Bureau-based industry measures of AI use linked to U.S. Bureau of Economic Analysis (BEA) and U.S. Bureau of Labor Statistics (BLS) production accounts. Their baseline results suggest that AI is productivity enhancing and input saving, although alternative timing assumptions yield weaker estimates and that progress in economic measurement remains essential for drawing stronger conclusions. Johnston and Makridis [2026] 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. [2026a] 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-

[2] OpenAI released the GPT-3 API in June 2020; by early 2021, over 300 applications were using it for customer service, writing, and coding, collectively generating 4.5 billion words per day [OpenAI and Pilipiszyn, 2021]. GitHub introduced Copilot as a technical preview in June 2021 and reported approximately 1.2 million preview users by its general release in June 2022. GitHub also reported that, in files for which Copilot was enabled, nearly 40 percent of code in supported popular programming languages was generated by Copilot [Friedman, 2021, Dohmke, 2022] These developments indicate that commercial experimentation and complementary investment preceded ChatGPT's public release. Enterprise writing assistants such as Jasper entered the market in 2021, and by late 2021 enterprise AI tools were embedded in production workflows across sectors. VC/ML investment in startups rose by 139% in 2021, evidence that enterprise diffusion and investment accelerated well before the ChatGPT became salient [Johnston and Makridis, 2026].

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), the Gallup measures exceed the corresponding BTOS measure by roughly 15 to 32 percentage points over overlapping periods [Bonney et al., 2024]. Gallup estimates the share of workers reporting personal AI use or organizational integration, whereas BTOS estimates the share of employer businesses reporting AI use in business functions. The Gallup measures are therefore worker-based, while BTOS is defined over firms. After BTOS broadened its AI question in late 2025, its aggregate series moved closer to Gallup's frequent-use series. Understanding the extent to which industries and workers are using new technologies like AI is essential for correctly measuring important macroeconomic outcomes like productivity. 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. Nonetheless, we document a high cross-industry correlation between the two measures, suggesting that they track a common diffusion process, but their levels are not directly comparable. Because our empirical strategy relies on relative variation in AI use across state-industry cells rather than on the national level of use, these measurement differences affect the estimates primarily insofar as they alter cross-cell rankings or introduce differential measurement error.[3]

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 time-invariant state$\times$industry and industry$\times$year shocks. We find that a 1 percentage point increase in the AI share of users in a state$\times$industry is associated with roughly a 0.4–0.6% increase in employment and a 0.1–0.2% increase in real output after 2020, during the early diffusion of enterprise AI tools.[4] The employment estimates are statistically noisy, whereas the real output estimates are more precisely estimated. As time elapses, exposure and realized use are likely to diverge, but the magnitudes are economically consequential without being implausible once mapped onto the observed pace of diffusion in AI use. For scale, an increase in the share of frequent AI users from roughly 12 percent to 24 percent would imply

[3] The distinction is particularly relevant for continuous-intensity designs. By contrast, Highfill and Samuels [2026a] classify industries into broad AI-intensity groups and therefore abstract from continuous differences in adoption intensity.

[4] Additional research is required to understand whether these results generalize in settings with greater quasi-experimental variation. For example, if telemetry data could be used to impute historical AI utilization at the state$\times$industry rates beyond those observed in our sample, then a strategy similar to Johnston and Makridis [2026] could be employed.

approximately 1.2 to 2.5 percent higher real output under our estimates.[5] Put differently, the implied magnitude is large relative to recent aggregate growth, even though the estimates should be interpreted as reduced-form state-by-industry associations rather than a national accounting decomposition of GDP growth. Relative to recent macroeconomic benchmarks, this places our reduced form estimates above the more conservative aggregate calibration in Acemoglu [2025], who projects only modest GDP gains over a 10-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.

An emerging empirical literature provides a more direct benchmark for these magnitudes. Bick et al. [2026b] relate worker-reported adoption to industry productivity and find that, in the United States, a 10 percentage point increase in AI adoption is associated with approximately 0.9 percentage points faster annualized labor-productivity growth from 2022:Q4 to 2025:Q3. They obtain comparable positive associations using firm adoption in Europe, but find no clear industry-level employment response. However, these relationships are also associations rather than causal effects, and their null employment result is not inconsistent with our imprecisely estimated positive employment coefficients. The sign and timing of our estimates also mirror Johnston and Makridis [2026], who construct state$\times$industry AI exposure by interacting task-level exposure with pre-2020 occupational composition in a difference-in-differences design. This predetermined treatment measure offers a more plausibly exogenous source of variation than contemporaneous realized use, although it identifies the effect of exposure rather than adoption. They find that a one-standard-deviation increase in AI exposure raises output by approximately 7 percent, with effects emerging in 2021; employment rises by roughly 4 percent where AI is more likely to augment workers and does not change significantly where it is more likely to automate tasks.

Firm-level evidence is likewise positive, but heterogeneous and often delayed. Using the adoption rates of matched U.S. peers as an instrument for adoption among European firms, Aldasoro et al. [2026] estimate a 4 percent increase in labor productivity, no adverse short-run employment effect, and larger gains among medium and large firms undertaking complementary investments in software, data, and workforce training. Based on a survey of nearly 750 corporate executives, Baslandze et al. [2026] infer annual revenue-based labor-productivity gains of roughly 0.4 percent in manufacturing, construction, and lower-skill services and 0.8 percent in finance in 2025, with larger gains anticipated in 2026. Most of these gains reflect revenue-based total factor productivity (TFP) and innovation- or demand-oriented

[5] 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, the rise from 11.8 percent in Q2:2024 to 24.2 percent in Q4:2025 documented in Makridis [2026] implies approximately 1.25 to 2.49 percent higher real output, reported in the text as 1.2 to 2.5 percent. Spread over the diffusion window from Q2:2024 to Q4:2025, this corresponds to about 0.8 to 1.7 percentage points of annualized real output growth. By comparison, real U.S. gross domestic product (GDP) grew 2.8 percent in 2024 and 2.1 percent in 2025, while real GDP rose by about 3.3 percent in total from Q2:2024 to Q4:2025, or roughly 2.2 percent at an annual rate. Thus, taken literally as an aggregate scale comparison, the implied AI-associated output gain would be equivalent to approximately 38 to 76 percent of actual real GDP growth over the same diffusion window. This comparison is intended to benchmark the magnitude of the estimate, not to attribute that share of aggregate GDP growth causally to AI.

channels rather than capital deepening. A more cautious picture emerges from Yotzov et al. [2026]: more than 80 percent of nearly 6,000 surveyed firms report no productivity or employment effect over the preceding 3 years, although executives expect AI over the next 3 years to raise productivity by 1.4 percent and output by 0.8 percent while reducing employment by 0.7 percent. Finally, aggregating projected sector-specific effects across 65 U.S. industries in a multisector general-equilibrium framework, Filippucci et al. [2026] conclude that AI could contribute up to 0.9 percentage points to annual aggregate TFP growth over the next decade. In this sense, the lower end of our implied annualized output effect is close to both the U.S. industry estimate in Bick et al. [2026b] and the upper-end aggregate benchmark in Filippucci et al. [2026], whereas the upper end of our range is more expansionary. These comparisons are necessarily approximate because reported time savings, realized use, occupational exposure, labor productivity, revenue-based TFP, and real output are distinct estimands.

This paper nonetheless faces several limitations, each of which we address directly. First, and most importantly, our AI use shares are measured in 2025 and Q1:2026 as a proxy of exposure corresponding to annual outcomes spanning 2016–2024. Because adoption is measured *after* the outcome window ends, reverse causality is a genuine concern: state-industry cells that experienced stronger output or productivity growth may have subsequently invested more in AI, generating a positive association between 2025 adoption rates and earlier economic performance that need not reflect AI's contribution to that performance. To address this, we instrument realized AI use with a predetermined measure of generative AI exposure constructed from occupation-level large language model (LLM) exposure scores [Eloundou et al., 2024] aggregated using 2015–2019 American Community Survey occupation shares, following Johnston and Makridis [2026]. Because these shares predate the diffusion of enterprise AI tools, the instrument is anchored in preexisting task composition rather than post-treatment adjustment. The instrumental variable (IV) point estimates for earnings, output, and productivity turn positive after 2020, but the confidence intervals are wide and none of the post-2020 coefficients is statistically distinguishable from zero; the employment estimates are similarly noisy. Second, the Gallup Workforce Panel is limited by sample size at the state-industry level, which forces us to aggregate across quarters to construct sufficiently populated cells, leaves many state-industry cells below Gallup's standard reporting threshold, and constrains our ability to measure within-cell changes in adoption over time. We address this by weighting observations by the number of respondents underlying each cell's AI measure and by re-estimating the level-based baseline specification on the subsample of cells with at least 30 respondents. The positive employment and real output patterns remain, although employment remains imprecisely estimated; the TFP point estimates are positive but are no longer statistically distinguishable from zero. Third, even with state$\times$industry and industry$\times$year fixed effects, identification remains challenging because AI adoption is not randomly assigned and may be correlated with managerial quality, complementary intangible capital, or local demand conditions that we do not fully observe. We also verify that our conclusions are robust to alternative normalizations, including using 2019 rather than 2020 as the reference year, which yields similar or somewhat stronger effects. For all these reasons, we interpret our estimates as important associations between productivity and actual AI use at the state-by-industry level, paving the way for future macroeconomic and econometric modeling of their causal effects.

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 Utilization Measures

We use two complementary measures of AI exposure—specifically, utilization. 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 (and "I don't know"). In our coding, "frequent" AI users are those who use AI daily or use a few times a week; those who respond "I don't know" are treated as infrequent, non-AI users.[6] 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 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 two-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.[7] The Census Bureau 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

[6] We have 52,729 observations for Q2–Q4:2025 and 22,563 observations for Q1:2026. For more details: https://www.gallup.com/699797/indicator-artificial-intelligence.aspx

[7] 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.

coarse, industry-only exposure design based on higher frequency business-reported AI use.

2.2. LEHD Employment and Earnings Data

To measure labor-market outcomes in the state-by-industry analysis, we use the Census Bureau's Quarterly Workforce Indicators (QWI), produced through the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI are observed quarterly and count worker–employer job matches rather than unique employed persons. Our baseline employment variable is beginning-of-quarter employment (`Emp`), which measures the number of jobs present on the first day of the reference quarter. For each state×two-digit-NAICS×year cell, we construct annual employment as the arithmetic mean of the available quarterly `Emp` observations and take its logarithm. A robustness exercise takes full-quarter employment (`EmpS`), restricting employment to jobs held with the same employer on both the first and last day of the reference quarter and therefore capturing a more stable set of job matches. We use `Emp` as the baseline because it more closely approximates the contemporaneous stock of jobs.

Our earnings outcome is average monthly earnings among full-quarter jobs (`EarnS`), rather than aggregate payroll. We construct the annual measure as the arithmetic mean of the available quarterly `EarnS` observations within each state-industry-year cell, deflate it using the annual personal consumption expenditures price index, and take its logarithm. Accordingly, the earnings estimates describe differences in the average real monthly earnings of workers in stable jobs, not differences in total state-industry earnings. We harmonize the QWI industry classifications to the two-digit NAICS categories.

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[8], allowing us to examine whether greater AI exposure is associated not only with labor market changes but also with changes in production.[9] 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. This measure is particularly valuable because it captures the productivity-related dimension of AI adoption more directly than employment or earnings alone.

[8] Real output is constructed using BEA's annual GDP by industry by state data as nominal GDP (the SAGDP2 tables) deflated with the aggregate GDP price index. We also found nearly identical results for output and TFP when using the GDP series in 2017 chained prices.

[9] 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.

2.4. Sample Construction

Our estimation sample covers annual observations from 2016 through 2024.[10] The outcome panel contains annual state-industry-year observations from 2016 through 2024. For each state-industry cell, we pool Gallup responses from Q2:2025 through Q1:2026 to construct a single survey-weighted frequent-use share, which is merged to every year of the outcome panel.[11] The primary analyses use the full individual-level Workplace Panel sample consisting of individuals of at least 18 years old and full and part -time workers. We then create weighted state-industry averages, which we use as model-based contextual predictors. However, because many state-industry cells fall below Gallup's standard reporting threshold, state-industry cells are not reported or interpreted as standalone Gallup estimates. Moreover, we also estimate robustness models restricted to cells with at least 100 respondents; conclusions are substantively similar, though less precise because this restriction requires discarding over two-thirds of the sample. 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 estimated 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

[10] The state×industry data is not yet available for 2025.

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

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 could 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 rises from roughly 7 to 10 percent, 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.[12] 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 in the most recent matched Q4:2025–Q1:2026 observations. 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 across business functions, 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. A shared disadvantage of the Gallup data is that sample size constraints mean it may not be fully representative of state-by-industry cells, but the sample continues to grow larger each quarter.

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

[12]Worker reports may capture informal use of personal AI tools at organizations that have not formally adopted, provided, or governed those tools. This is another reason worker- and employer-reported measures need not coincide.

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 LLMs, 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 output during the early diffusion of enterprise generative AI tools. While the empirical strategy will not allow us to make causal statements, the results will tell us whether the states and industries that exhibit higher shares of their workers using AI actively also display greater employment and productivity growth over the duration of the sample, which contributes to an ongoing policy debate on productivity and AI.

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. (Results are robust to a 2019 normalization too.)

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 purges variation that might be coming from industry-specific booms and busts, e.g. the technology sector during the pandemic. Similarly, the inclusion of α_{si} purges variation that might be coming from longstanding state-by-industry human capital differences, like California's technology sector or New York's finance sector. 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 utilization over Q2:2025–Q1:2026, 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 retrospective 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 similar in spirit to Highfill and Samuels [2026a], who interact an industry-level measure of AI intensity with a post-2021 indicator in a difference-in-differences specification containing industry and year fixed effects.

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. As more data become available, this question can be answered.

Beyond reverse causality, the empirical setup is well-suited to address the reality 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. 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 empirical specification with the Gallup data especially demanding, since identification comes within-state-by-industry changes in outcomes over time, with those changes compared across states within the same industry according to differences in AI use.

The identifying assumption for the event study is borne out in parallel trends: absent differences in AI utilization, 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 pretrends 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$\times$industry and industry$\times$year, which are critical for addressing a suite of potential omitted variables.

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, 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.[13]

Panel B shows a cleaner pattern for real output, consistent with Johnston and Makridis [2026]. 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 employment estimates do not provide precise evidence of either displacement or expansion, while the real-output estimates show that cells with greater subsequent AI use experienced stronger post-2020 output paths.

A note of caution is warranted when comparing the relative magnitudes of these estimates. Notably, the estimated employment semi-elasticity is larger than the corresponding real-output semi-elasticity. If the coefficients are interpreted jointly, this ordering implies lower rather than higher output per worker: state-industry cells with a 1 percentage point higher share of frequent AI users exhibit approximately 0.4–0.6 percent higher employment but only 0.1–0.2 percent higher real output. The difference between the output and employment coefficients therefore corresponds mechanically to a negative association with labor productivity. This comparison should not be overinterpreted, however, because the employment estimates are imprecise, and the difference between the two coefficients may not be statistically distinguishable from zero. The pattern could also reflect an early stage of adjustment in which firms expand employment or hire complementary workers before the associated output gains are fully realized. More generally, the estimates identify reduced-form associations rather than structural production-function parameters and cannot distinguish among labor hoarding, organizational adjustment, demand expansion, or measurement differences across the outcomes. The employment pattern is similar when we instead use the Quarterly Workforce Indicators (QWI) measure of stable, full-quarter employment: the post-2020 coefficients remain positive, are largest in 2022 and 2023, but move closer to zero in 2024.[14]

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

[14] Our baseline outcome is beginning-of-quarter employment (`Emp`), which counts worker–employer job matches observed in the preceding and reference quarters and is intended to approximate the stock of jobs on the first day of the reference quarter. The alternative full-quarter employment measure (`EmpS`) additionally requires the same job match to be observed in the subsequent quarter, thereby restricting the outcome to jobs that persist throughout the reference quarter. We prefer `Emp` because it is the QWI employment concept closest to conventional QCEW employment and because our object of interest is the broader employment stock rather than only jobs that survive into the following quarter. Conditioning on subsequent-quarter survival may attenuate changes that operate through job separations or changes in job duration.

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 $\times$ industry) and log private fixed assets (only industry). We find similar results.[15,16] In contrast, Figure 8 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 industry-state cells from the Workforce Panel could lead to measurement error, Figure 9 replicates the baseline after restricting to cells with at least 30 respondents determining the share of frequent AI users. The post-2020 employment and real output point estimates remain positive. Employment remains imprecisely estimated, while the positive real output pattern is preserved; the TFP estimates are positive but are no longer statistically distinguishable from zero.

These estimates 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 11.8 percent in Q2:2024 to 24.2 percent in Q4:2025 would correspond to approximately 1.2 to 2.5 percent higher real output. Spread over that diffusion window, this back of the envelope calculation implies roughly 0.8 to 1.7 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 [2025] 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 [2026], 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 preexisting occupational composition rather than post treatment adjustments in employment or task allocation. We then interact this prepandemic 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 prepandemic

[15] We also examined whether growth in AI use at a state$\times$industry level predicts higher economic activity. The Gallup Workforce Panel is not yet large enough to create precise within state-industry growth rates, so many cells are too small to include. The post-2020 log employment point estimates are positive for states and industries with higher growth in AI use, but the displayed confidence intervals include zero; the evidence is suggestive rather than statistically conclusive.

[16] 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.

task mix, experienced larger post 2020 changes in outcomes through their induced increases in AI use. Identification comes from whether state-industry cells with greater predetermined exposure experience larger changes over time than less-exposed cells in the same industry, after absorbing time-invariant differences across state-industry cells and annual shocks common to an industry across states.

Figure 10 shows positive post-2020 IV point estimates for earnings, real output, and productivity, but the estimates are highly imprecise and none of the post-2020 coefficients is statistically distinguishable from zero. Panel A shows that the coefficients for log employment are noisy throughout the sample and do not display the same post-2020 pattern as in the baseline estimates. The point estimates are positive in several years, but the confidence intervals are wide and include economically meaningful negative effects as well, which suggests that the employment response is not robust to using preperiod generative AI exposure as an instrument for realized AI use.

The remaining panels provide directionally supportive, but not statistically conclusive, evidence. Panel B shows positive post-2020 point estimates for log earnings. Panel C displays the strongest positive point-estimate pattern for real output, with coefficients rising after 2020 and remaining positive through the end of the sample. Panel D shows a similar positive point-estimate pattern for total factor productivity. Because all of the displayed post-2020 confidence intervals include zero, these estimates cannot establish that the output pattern reflects efficiency gains rather than input accumulation. The IV exercise is therefore best interpreted as a directional robustness check rather than independent statistical confirmation of the baseline estimates; the first-stage simply is not strong enough.

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 research 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 Business School [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 11 examines whether higher AI use is associated with shifts in the composition of intangible capital across industries. The evidence is mixed rather than uniformly 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 by 2021 and turn positive thereafter; the 2022 and 2023 estimates are positive and their displayed pointwise 95 percent confidence intervals exclude zero. Software shares (Panel C) follow a mixed and imprecisely estimated trajectory—positive early in the sample, declining through 2021–2022, and turning negative by 2023. Data capital shares (Panel D) are mostly negative, with the largest negative point estimates in 2016–2017 and an estimate close to zero in 2022 before turning negative again in 2023.[17] Taken together, the four panels do not show broad, portfolio-wide reallocation of intangible capital, but they do show a localized positive R&D association after 2020. Because the analysis examines multiple exploratory outcomes without a multiple-testing adjustment, the isolated R&D estimates should be interpreted cautiously. AI use at the industry level may also not yet register broadly in official capital account data—consistent with the 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 conclusive evidence of AI-driven capital reallocation. The absence of broad effects could also reflect the lack of within-industry geographic variation that sharpens identification in the baseline employment and output results.

We present these results with the measurement framework by Corrado et al. [2005] and Corrado et al. [2009]. The official national accounts that we draw upon include a large portion of 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. Only the R&D share displays statistically positive post-2020 estimates in 2022 and 2023 under the displayed pointwise confidence

[17] 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].

intervals; the other outcomes do not show robust post-2020 responses. This pattern 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 broad signal in our industry-level specifications is therefore not strong evidence against AI-driven capital deepening; it may instead reflect the "asset boundary" of official statistics.

Corrado et al. [2025] 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 11, we have included data capital services based on the work in Samuels et al. [2026]. Based on those working estimates, however, AI intensity is not associated with a shift toward data assets—at least not yet.

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 efforts to improve and expand measurement 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 worker-reported AI use to administrative and national accounts data from LEHD and BEA over 2016 to 2024, rather than relying only on occupational exposure proxies or broad technological classifications. Exposure measures are informative about where AI could matter, whereas realized-use measures describe where workers report that AI use actually occurs—and AI use explains which production processes, organizational routines, and worker engagement 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 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 to combine telemetry-based measures of software interaction, platform activity, and workflow integration with representative 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 or states. 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 uneven, but already expanding, enterprise diffusion of AI before 2022.

References

Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, January 2025.

Philippe Aghion and Simon Bunel. AI and growth: Where do we stand? Policy note, June 2024. URL https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf.

Iñaki Aldasoro, Leonardo Gambacorta, Rozalia Pal, Debora Revoltella, Christoph Weiss, and Marcin Wolski. AI adoption, productivity and employment: Evidence from european firms. BIS Working Paper 1325, Bank for International Settlements, January 2026.

Lefteris Andreadis, Eleni Kalotychou, Manolis Chatzikonstantinou, Christodoulos Louca, and Christos A. Makridis. Local heterogeneity in artificial intelligence jobs over time and space. *AEA Papers and Proceedings*, 115:29–34, May 2025.

Salomé Baslandze, Zachary Edwards, John R. Graham, Ty McClure, Michael Sparks, Brent Meyer, Sonya Ravindranath Waddell, and Daniel Weitz. Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. Working Paper 2026-4, Federal Reserve Bank of Atlanta, March 2026.

Alexander Bick, Adam Blandin, and David J. Deming. The rapid adoption of generative AI. *Management Science*, January 2026a. doi: 10.1287/mnsc.2025.02523. Articles in Advance; published online January 20, 2026.

Alexander Bick, Adam Blandin, David J. Deming, Nicola Fuchs-Schündeln, and Jonas Jessen. Mind the gap: AI adoption in europe and the U.S. Brookings Papers on Economic Activity, Spring 2026 Conference Draft. Also circulated as NBER Working Paper No. 34995, March 2026b.

Kathryn Bonney, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, and Keith Savage. 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, 2024.

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

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab, November 2025. Revised November 13, 2025.

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

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

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

Carol Corrado, Jonathan Haskel, Cecilia Jona-Lasinio, and Massimiliano Iommi. Intangible capital and modern economies. *Journal of Economic Perspectives*, 36(3):3–28, 2022.

Carol Corrado, Jonathan Haskel, Massimiliano Iommi, Cecilia Jona-Lasinio, and Filippo Bontadini. Data, intangible capital, and productivity. In Susanto Basu, Lucy P. Eldridge, John C. Haltiwanger, and Erich H. Strassner, editors, *Technology, Productivity, and Economic Growth*, volume 83 of *Studies in Income and Wealth*, pages 193–243. University of Chicago Press, Chicago, 2025.

Thomas Dohmke. GitHub Copilot is generally available to all developers. The GitHub Blog, June 2022. URL https://github.blog/news-insights/product-news/github-copilot-is-generally-available-to-all-developers/. Accessed July 19, 2026.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024.

Francesco Filippucci, Peter Gal, and Matthias Schief. Aggregate productivity gains from artificial intelligence: A sectoral perspective. *AEA Papers and Proceedings*, 116:31–35, May 2026. doi: 10.1257/pandp.20261035.

Nat Friedman. Introducing GitHub Copilot: Your AI pair programmer. The GitHub Blog, June 2021. URL https://github.blog/news-insights/product-news/introducing-github-copilot-ai-pair-programmer/. Accessed July 19, 2026.

Tina Highfill and Jon D. Samuels. Early estimates of the impact of ai within bea's industry economic accounts. BEA Working Paper WP2026-3, U.S. Bureau of Economic Analysis, February 2026a. BEA Working Paper Series.

Tina Highfill and Jon D. Samuels. 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, 2026b. Forthcoming, BEA Working Paper Series.

Tina Highfill, David Wasshausen, and Gregory Prunchak. 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, January 2025. BEA Working Paper Series.

Anders Humlum and Emilie Vestergaard. Still waters, rapid currents: Early labor market transformation under generative AI. NBER Working Paper 33777, National Bureau of Economic Research, May 2025. Revised March 2026; previously circulated as "Large Language Models, Small Labor Market Effects".

Andrew C. Johnston and Christos A. Makridis. AI, output, and employment. CESifo Working Paper 12579, CESifo, March 2026.

Christos A. Makridis. The organizational transmission of AI: The role of managers on AI adoption and impact. CESifo Working Paper 12373, CESifo, April 2026. Revised April 2026.

Christos A. Makridis and Erik Brynjolfsson. Counting AI: A blueprint to integrate AI investment and use data into US national statistics. Brookings Institution, January 2026. URL https://www.brookings.edu/articles/counting-ai-a-blueprint-to-integrate-ai-investment-and-use-data-into-us-national-statistics/. Accessed July 19, 2026.

Joshua New, Marina Meyjes, and Austin Carson. The u.s. needs a generative ai intensity index. Technical report, SeedAI, October 2025.

OpenAI and Ashley Pilipiszyn. GPT-3 powers the next generation of apps. OpenAI, March 2021. URL https://openai.com/index/gpt-3-apps/. Accessed July 19, 2026.

Jon D. Samuels, J. Bayoán Santiago Calderón, and Corby Garner. The impact of capitalizing data on productivity growth in the u.s. BEA Working Paper WP2026-7, U.S. Bureau of Economic Analysis, March 2026. URL https://www.bea.gov/system/files/papers/BEA-WP2026-7.pdf. BEA Working Paper Series.

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

Ivan Yotzov, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J. Davis, Kevin M. Foster, Aaron Jalca, Brent H. Meyer, Paul Mizen, Michael A. Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang. Firm data on AI. NBER Working Paper 34836, National Bureau of Economic Research, February 2026.

Tables and Figures

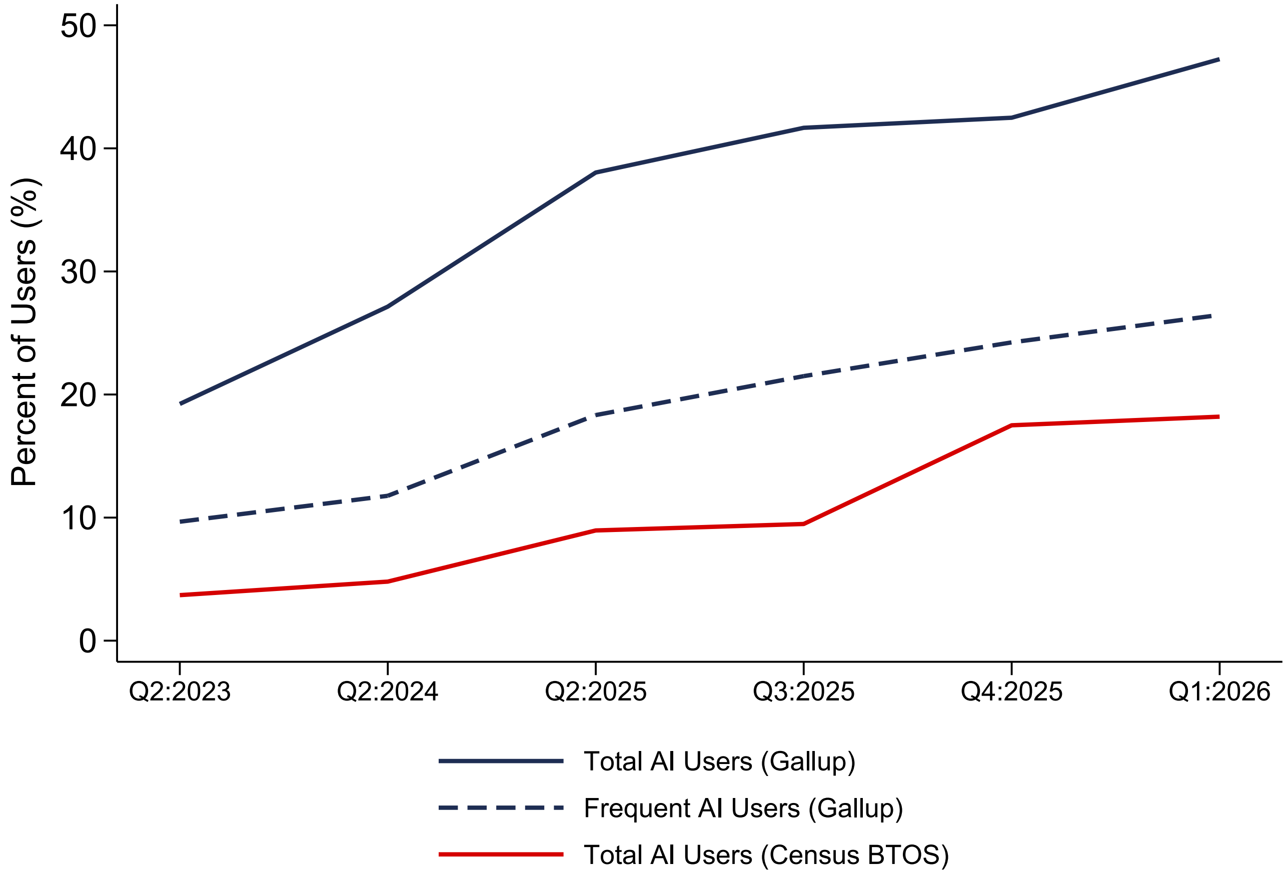

Figure 1. Comparing the Gallup and Census AI Measures Over Time

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.

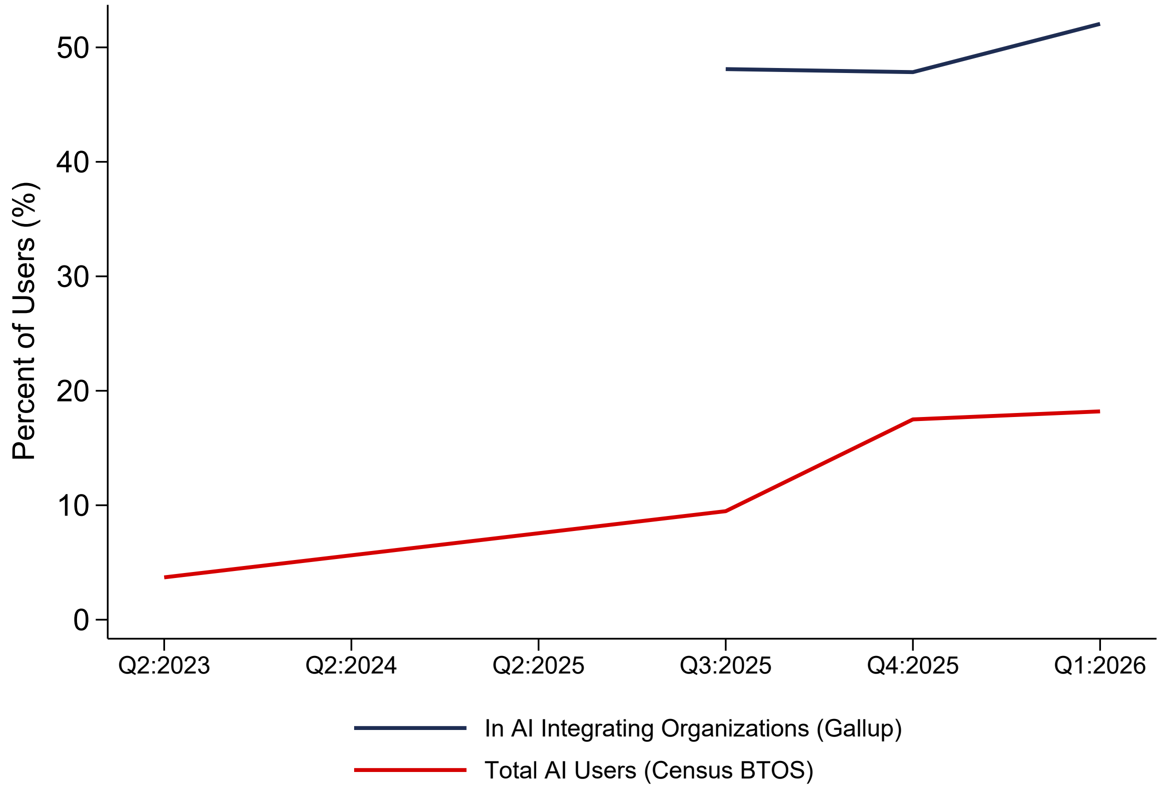

Figure 2. Worker Perceptions Versus Firm-Reported Adoption

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, I don't know) differ 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. (Beginning in Q3:2025, Gallup added a "Don't know" response, which is included in the denominator.)	"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+); annual prior	Quarterly (Q2:2025+)	Fortnightly, aggregated to quarterly
Sample size	∼20,000 workers per quarter, national	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 15–32 percentage points reflects genuine conceptual differences rather than simple undercounting by either survey.

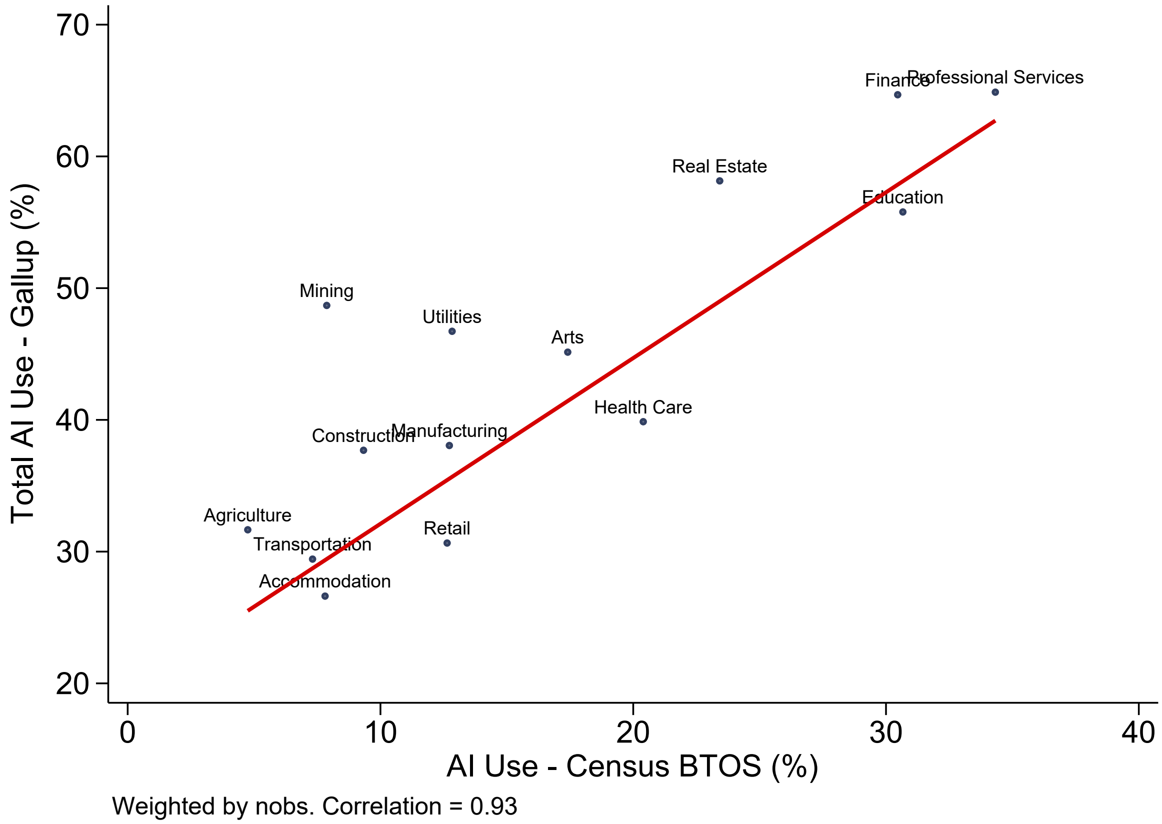

Figure 3. Comparing the Gallup and Census AI Measures Across Industries

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

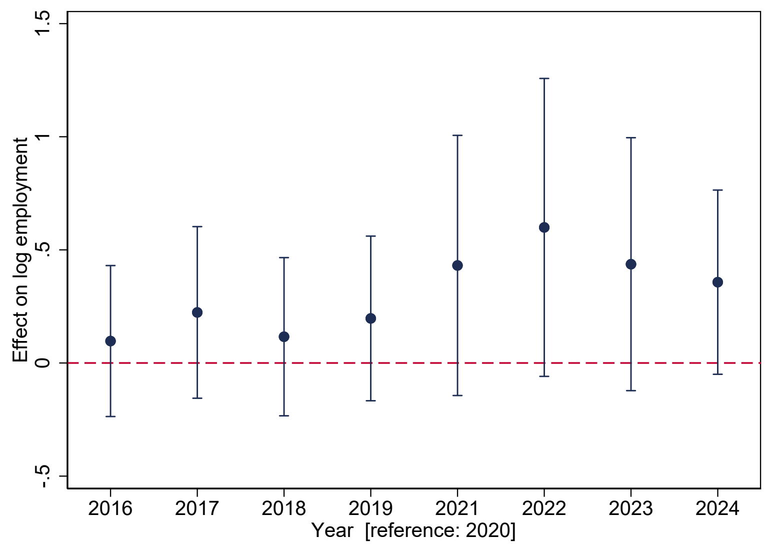

a. Employment

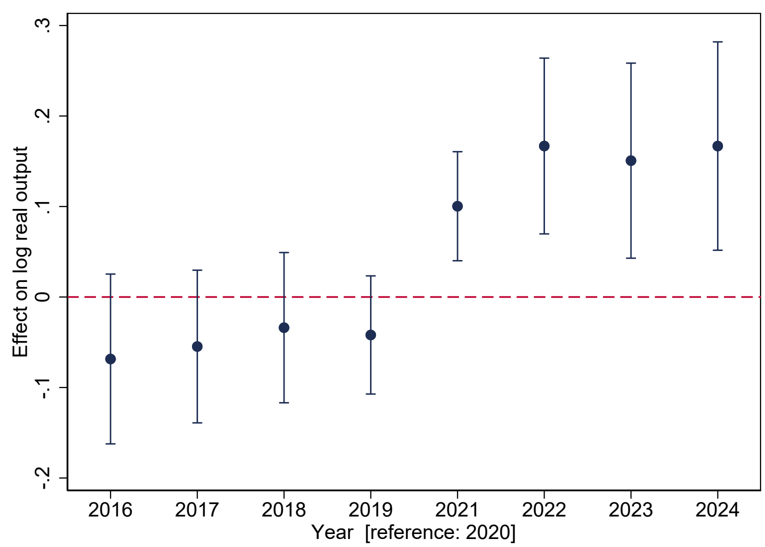

b. Real Output

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.

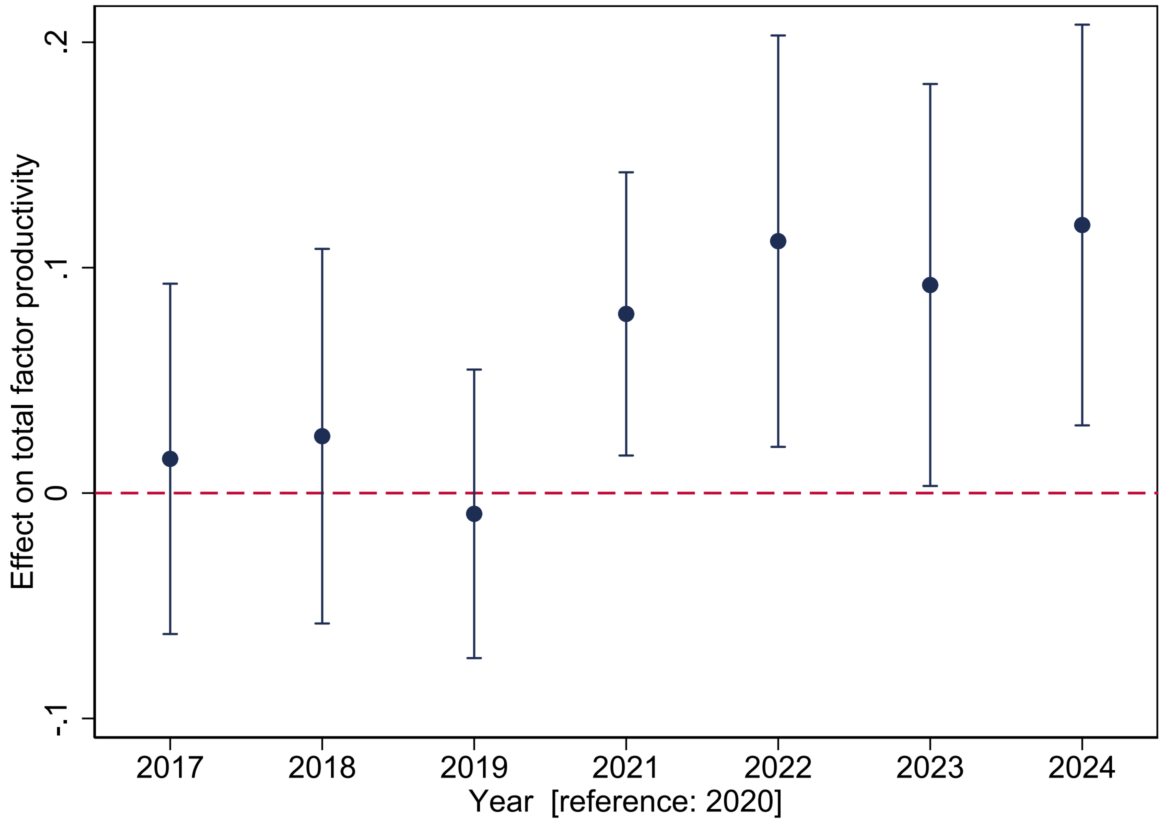

Figure 5. Event Study Robustness on Total Factor Productivity

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. Effect of AI Exposure on Log Employment, by Firm Size

a. Panel A: 0–19 employees

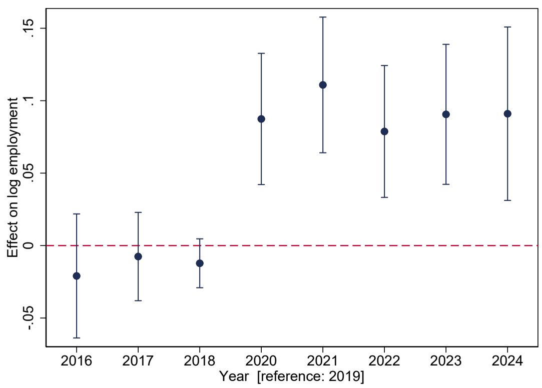

b. Panel B: 20–49 employees

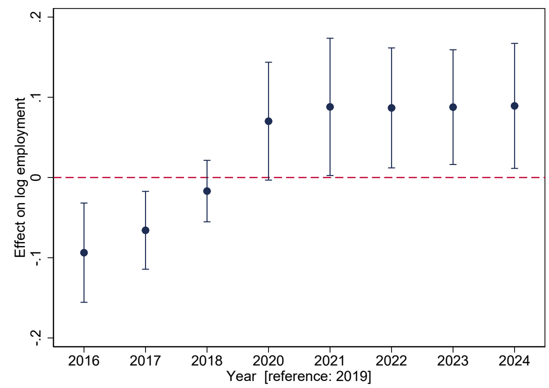

c. Panel C: 50–249 employees

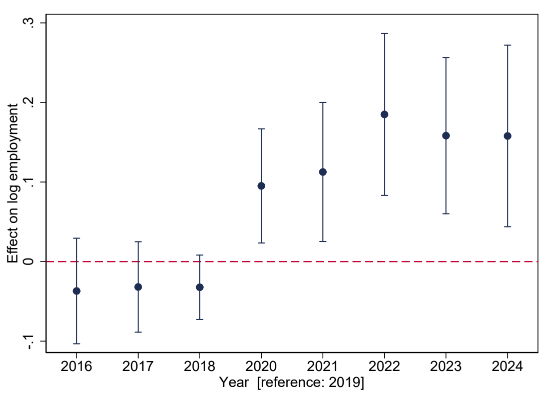

d. Panel D: 250+ employees

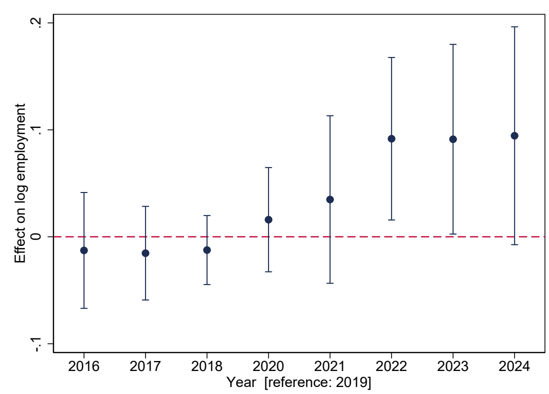

Notes.—Source: Gallup survey data and the Longitudinal Employer-Household Dynamics (LEHD) Quarterly Workforce Indicators, 2016–2024. The figure reports the coefficients associated with regressions of log employment on the Gallup-based AI share measure interacted with year fixed effects, conditional on state by industry and industry by year fixed effects, estimated separately by firm-size class. Employment is the annual mean of quarterly firm employment within each state by industry by firm-size cell. Panel D pools the 250–499 and 500+ size classes. Coefficients are normalized relative to 2019, the omitted reference year. Observations are weighted by each cell's baseline (2019) employment; the sample is restricted to cells with at least three nonmissing quarters. Standard errors are clustered at the state by industry level.

Figure 7. Effect of AI Exposure on Log Employment, by Firm Age

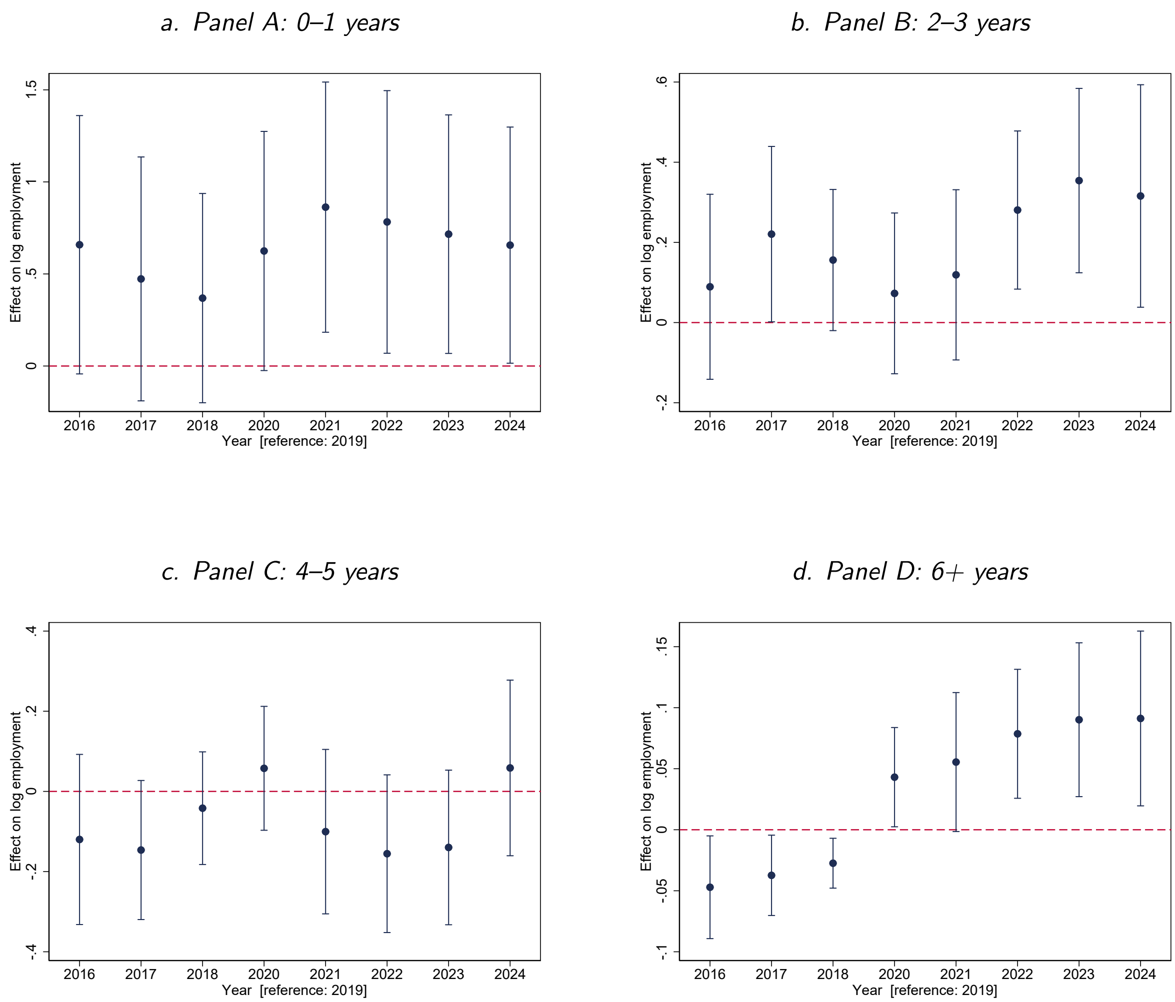

Notes.—Source: Gallup survey data and the Longitudinal Employer-Household Dynamics (LEHD) Quarterly Workforce Indicators, 2016–2024. The figure reports the coefficients associated with regressions of log employment on the Gallup-based AI share measure interacted with year fixed effects, conditional on state by industry and industry by year fixed effects, estimated separately by firm-age class. Employment is the annual mean of quarterly firm employment within each state by industry by firm-age cell. Panel D pools the 6–10 year and 11+ year age classes. Coefficients are normalized relative to 2019, the omitted reference year. Observations are weighted by each cell's baseline (2019) employment; the sample is restricted to cells with at least three nonmissing quarters. Standard errors are clustered at the state by industry level.

Figure 8. Baseline Event Study Results Using BTOS AI Shares

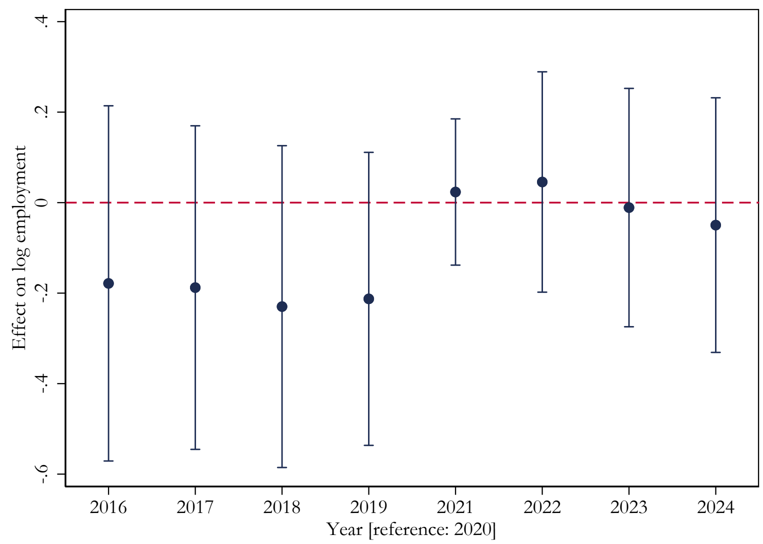

a. Employment

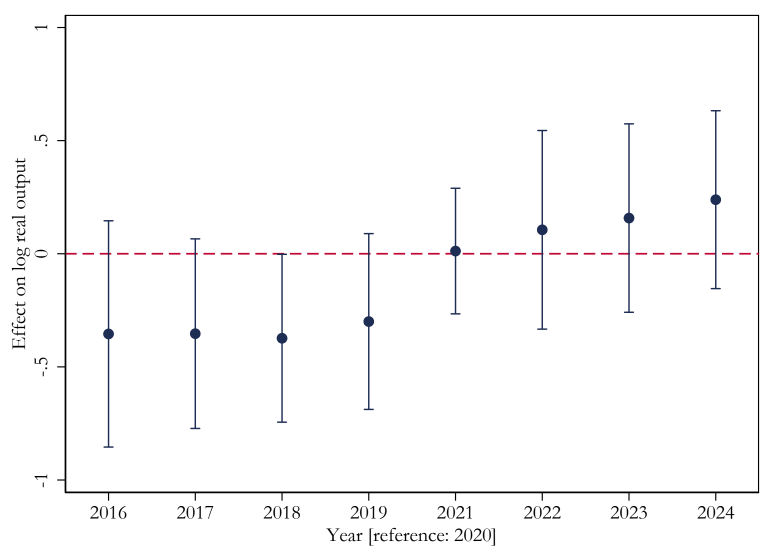

b. Real Output

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 9. Robustness Restricting Baseline to Larger State×Industry Cells

a. Panel A: log(Employment)

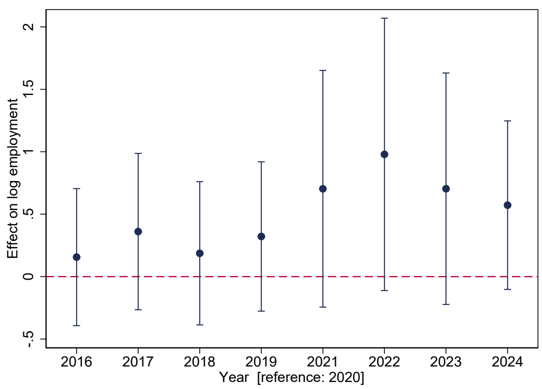

b. Panel B: log(Earnings)

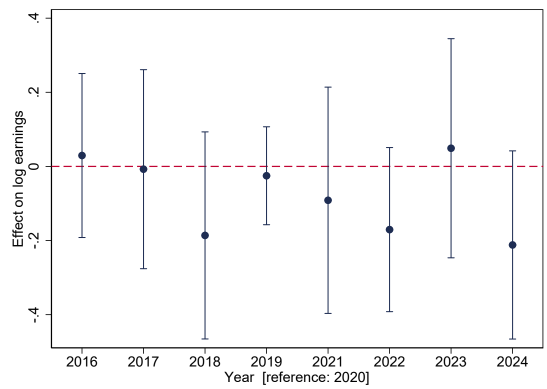

c. Panel C: log(Real Output)

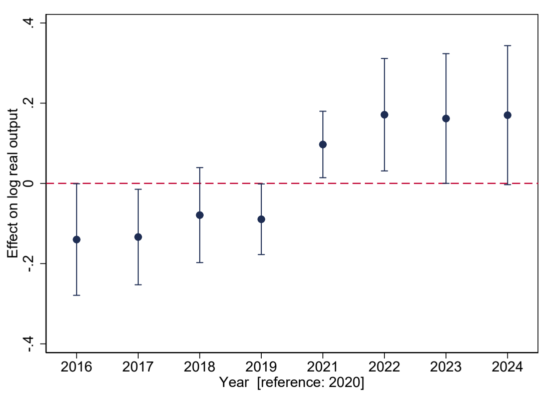

d. Panel D: Total Factor Productivity

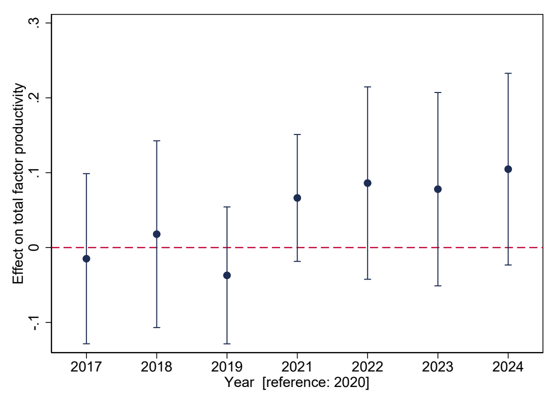

Notes.—Source: Gallup survey data, Longitudinal Employer-Household Dynamics (LEHD), and Bureau of Economic Analysis (BEA), 2016–2024. The figure reports the coefficients associated with regressions of log employment, log earnings, log real output, and total factor productivity (TFP) on the Gallup based 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 10. Robustness Using Generative AI Exposure to Instrument for AI Utilization

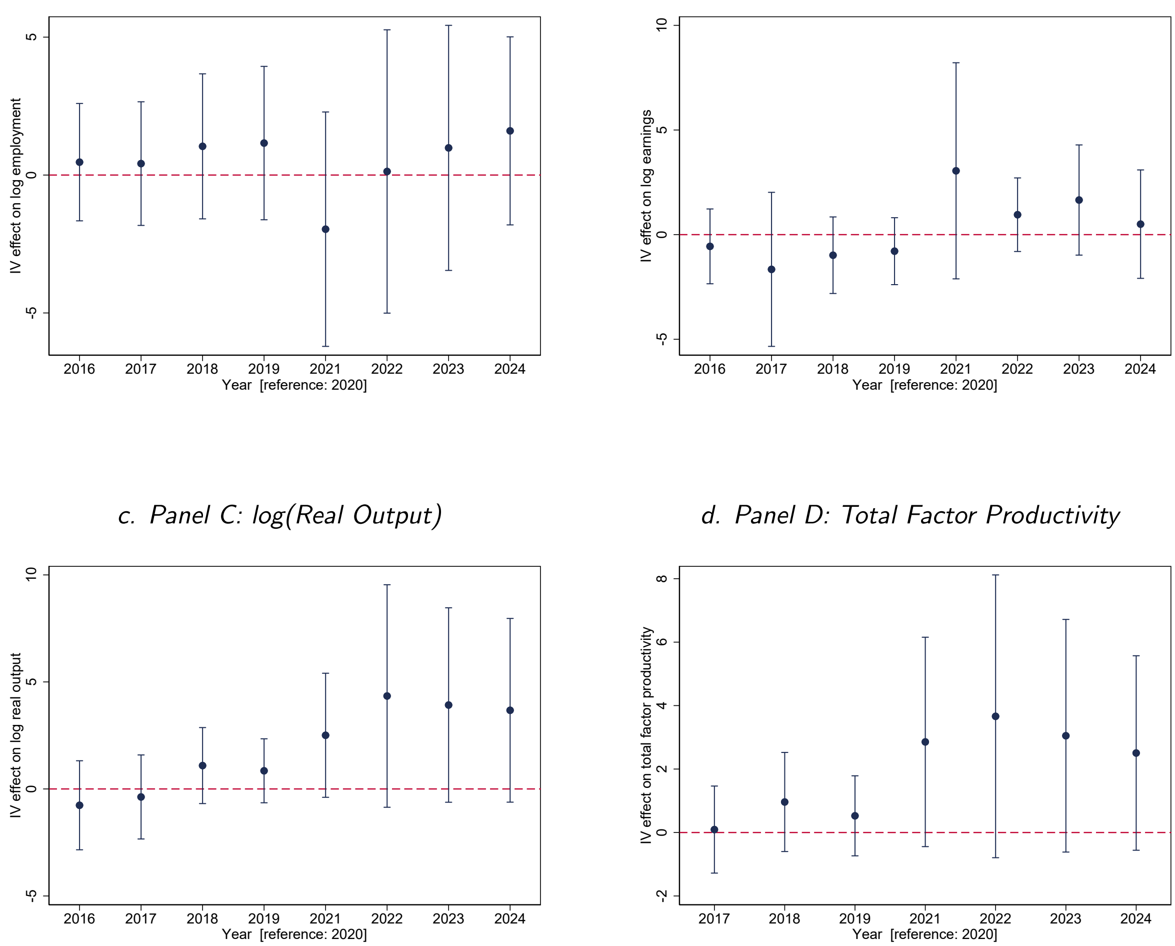

a. Panel A: log(Employment)

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 [2026], 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 11. Dynamic Effects of AI Use on Intangible Capital Shares

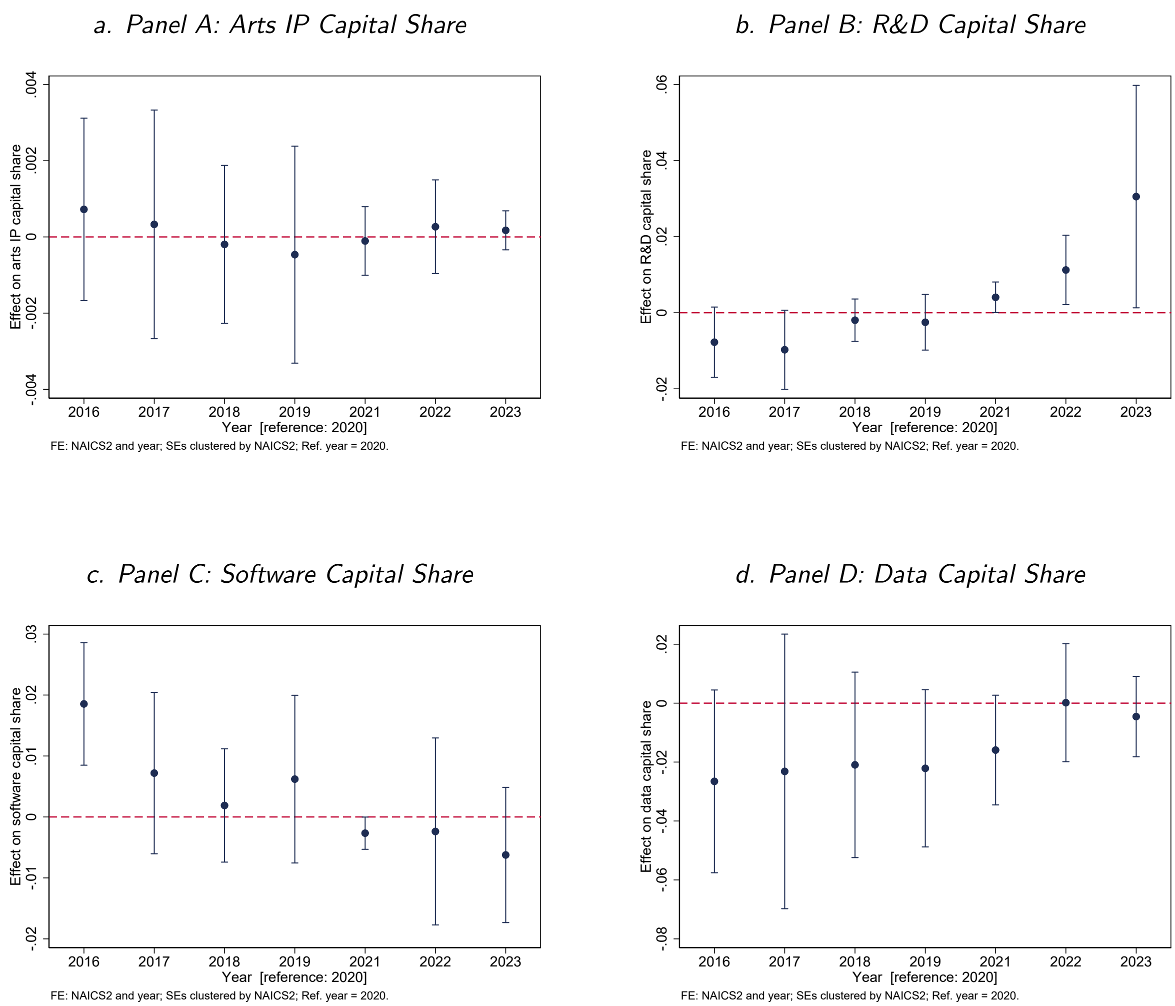

Notes.—Sources: Gallup Workforce Panel (2025); Integrated Labor Productivity Accounts (ILPA), 2016–2023. The figure reports coefficients from regressions of the industry share of arts IP, R&D, software, and data capital on the Gallup-based measure of frequent AI use, interacted with year indicators and 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.