

NBER WORKING PAPER SERIES

THE POLITICS OF AI

Nicholas Bloom
Christos Makridis

Working Paper 34813
<http://www.nber.org/papers/w34813>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026, Revised March 2026

We would like to thank the team at Gallup for their review of the data and draft. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research or any affiliated institutions.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34813>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Nicholas Bloom and Christos Makridis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Politics of AI
Nicholas Bloom and Christos Makridis
NBER Working Paper No. 34813
February 2026, Revised March 2026
JEL No. J0

ABSTRACT

Using new data from the Gallup Workforce Panel, we show that the apparent partisan divide in workplace AI adoption is largely an artifact of educational and occupational sorting rather than ideological differences in technology adoption. While Democrats are consistently more likely than Republicans to report frequent AI use—30.1% versus 25% in Q1:2026—and exhibit deeper task-level integration across a broader range of work activities, this raw gap shrinks to statistical insignificance once we control for educational attainment, and reverses sign in fully saturated models with occupation and industry fixed effects. Democrats work in more AI-exposed occupations and more AI-supportive organizational environments, but these differences reflect where people work and what skills they hold, not political identity per se. The one dimension that resists this compositional explanation is perceived job displacement risk: partisan differences were statistically indistinguishable through 2025 but a 4.2 percentage point gap has emerged by Q1:2026, suggesting that belief divergence may be beginning to outpace the structural factors that explain behavioral differences. In this sense, the “politics of AI” is better understood as the political geography of human capital and occupational structure than as an ideological cleavage in technology adoption.

Nicholas Bloom
Stanford University
Department of Economics
and NBER
nbloom@stanford.edu

Christos Makridis
Arizona State University
and Gallup and Stanford University
cmakridi@asu.edu

1 Introduction

Large language model (LLM) technologies, such as OpenAI’s ChatGPT and Anthropic’s Claude, have diffused at an unprecedented pace, generating surging economic and policy interest (Figure 1). Since the public release of ChatGPT in late 2022, generative AI tools have attracted hundreds of millions of users and were being queried billions of times per month by early 2024 (Bick et al., 2025; Makridis, 2025). Researchers have quickly recognized that these models could significantly transform work, given their ability to perform a broad range of cognitive tasks. Indeed, Eloundou et al. (2024) found that around 80% of the U.S. workforce could have at least 10% of their job tasks affected by LLMs, and nearly one-fifth of workers might see 50% or more of their tasks impacted. Yotzov et al. (2026) suggest the adoption of AI could increase U.S. productivity growth by 0.7% a year from 2026, or potentially more as a general-purpose technology (Brynjolfsson et al., 2021). This potential, combined with the rapid improvements in model capabilities, has spurred concern and enthusiasm in equal measure—and raised the question of whether AI’s benefits and risks are distributed evenly across the workforce, including across political lines.

Not all occupations or workers are equally exposed to generative AI’s capabilities, creating uneven welfare effects. Researchers have developed occupation-level exposure measures by aligning job task descriptions with AI and LLM competencies (Felten et al., 2021; Eloundou et al., 2024).¹

In contrast to what we know about automation exposure,² these analyses suggest that higher-

¹An alternative approach has involved using AI job postings at a county-level (Andreadis et al., 2025) and firm-level (Acemoglu et al., 2022; Babina et al., 2024). These approaches suggest similar heterogeneity, but differ in that they focus more on AI-specific jobs, rather than the more general use of genAI by a wider set of employees.

²Much of the prior literature on automation exposure focuses on the substitution of routine tasks, with the highest exposure concentrated in middle-skill occupations and contributing to job polarization. Autor et al. (2003) formalize this task-based view, while Acemoglu and Restrepo (2018) and Acemoglu and Restrepo (2020) generalize it to a broader framework in which new technologies reallocate tasks between labor and capital, often with displacement effects concentrated outside high-skill cognitive work.

wage, highly educated occupations have especially high exposure to generative AI. Early adoption patterns line up with this exposure heterogeneity. Although estimates vary based on sampling methods, between 20-30% of U.S. employed workers reported actively using generative AI in their job, but adoption has been far from uniform. Usage rates are twice as high among workers with a college degree compared to those without, and considerably lower for older workers (ages 50+) relative to younger cohorts. Generative AI use is most common in professional and tech-oriented roles, whereas many lower-skill or manual occupations have seen little direct AI uptake so far. These early disparities raise the possibility that generative AI could widen existing skill and income gaps, as those workers most equipped to leverage the technology reap the initial benefits.

Emerging research on the labor market impacts of generative AI, especially from randomized and quasi-experimental settings, underscores substantial heterogeneity in who benefits from these tools. In controlled experiments with white-collar workers, access to LLM-based assistants delivers large average productivity gains but disproportionately helps lower-skilled or less experienced workers. [Noy and Zhang \(2023\)](#) randomize access to ChatGPT for mid-skill professionals performing writing tasks and find roughly 40 percent reductions in completion time and sizable improvements in output quality, with the largest gains at the bottom of the initial performance distribution. Similarly, [Brynjolfsson et al. \(2025b\)](#) study the introduction of a GPT-based assistant in a large U.S. customer-support operation and document double-digit productivity increases concentrated among novice agents, with much smaller effects for experienced workers. Among other similar RCTs, growing evidence portrays generative AI as a skill-equalizing technology within firms, narrowing performance gaps by raising the floor more than the ceiling—even as the underlying tasks remain highly complementary to cognitive and communication skills.

At a more aggregate level, however, macro evidence paints a more nuanced picture of how these

micro-productivity effects translate into wages and employment. Using high-frequency payroll data, [Brynjolfsson et al. \(2025a\)](#) document a pronounced relative decline in employment for 22-25 year-olds in the most AI-exposed occupations, suggesting that adjustment operates partly through reduced hiring or separations for young workers in roles where AI is more substitutable than complementary. However, the picture is not uniformly negative: leveraging linked Danish survey and administrative data, [Humlum and Vestergaard \(2025\)](#) find precise near-zero effects of chatbot adoption on earnings and hours over the first two years of diffusion, alongside evidence of task reshuffling and occupational mobility rather than changes in total labor input, whereas [Johnston and Makridis \(2025\)](#) find more positive effects on employment, wages, and output through 2024 arising from reallocation across occupations. These macro studies indicate that generative AI’s labor-market impact is highly context dependent, varying by country, institutional environment, exposure margin, and career stage with sizable productivity effects at the task level but aggregate wage and employment responses that can be positive, muted, or even adverse for specific groups.

One dimension of heterogeneity that has received little attention thus far is political affiliation. This is an important omission: political ideology can shape perceptions of and interactions with new technologies. In the United States, technology adoption and trust in science often diverge along partisan lines ([Mills and St. Clair, 2025](#)). For instance, Republicans express significantly lower trust in scientists and are less likely to embrace certain innovations than their Democratic counterparts. Early evidence suggests that partisan attitudes extend to AI as well. Recent surveys indicate that Americans in both parties perceive popular AI systems as having a political bias—specifically a leftward slant—which could influence their willingness to use such systems ([Hall et al., 2025](#)). Given these patterns, it is plausible that partisan identity might be an independent predictor of generative AI adoption: workers with different political leanings could have different

levels of exposure to AI-related information, distinct workplace cultures, or varying trust in AI, all of which might affect their propensity to use these new tools.

In this paper, we use new Gallup Workforce Panel data from Q2:2024-Q1:2026 to study partisan differences in AI adoption in the U.S. labor market. The Gallup Workforce Panel (GWP) is a nationally representative longitudinal survey of U.S. adults, with roughly 20,000 respondents per wave in recent years. The panel’s large sample and rich set of employment variables make it ideal for measuring technology usage across diverse worker groups. We focus on the data through Q1:2026 survey waves, which included questions on not only the frequency of broad AI utilization on-the-job, but also the composition. We combine the survey data with external measures of generative AI exposure at the occupation level. In particular, we assign each respondent an LLM exposure score based on their Standard Occupational Classification (SOC) code, using the task-level exposure index from [Eloundou et al. \(2024\)](#). This exposure metric proxies the technological opportunity or relevance of generative AI in the respondent’s job, controlling for other factors.

Our empirical approach is to compare AI usage rates between Republican-affiliated and Democratic-affiliated workers (allocating Independents based on past self-identified status), controlling for a host of confounding factors. In our baseline specifications, we include fixed effects for broad occupation (SOC major group) and industry, and a rich set of demographics (age, sex, race, education, marital status, and parent status). These controls ensure that we are effectively comparing Republicans and Democrats in similar job contexts—that is, a Democrat and a Republican of the same age, working in the same type of occupation with comparable AI exposure—to isolate any partisan gap in AI adoption. We estimate models for both an extensive margin (any AI use on the job) and an intensive margin (frequency of use) to assess how deep the differences run.

Our results document a persistent partisan gradient in workplace generative AI utilization, but

the source of the gradient is largely compositional. In the raw data, Democrats are more likely than Republicans to report frequent AI use at work, and the gap widens over 2024 to 2026. For example, the share reporting weekly or more frequent use rises from 13.2% (Democrats) versus 9.6% (Republicans) in Q2:2024 to 30.1% versus 25.0% in Q1:2026. Democrats also report greater breadth of use across work activities, as captured by an AI Productivity Index. In addition, partisan differences in perceived job displacement risk are modest and not precisely estimated.

However, once we condition on education and job characteristics, the partisan utilization gradient attenuates sharply. In the raw data, Democrats are 4.9 percentage points more likely to use AI frequently, but the differences attenuate after adding educational attainment, and even flip signs in fully saturated models with both occupation and industry fixed effects, implying Democrats are 2.2 percentage points less likely to use AI frequently. These results indicate that the raw partisan gap in AI adoption is not only much smaller than traditional accounts by the media, but also can be explained by differences in educational attainment and the distribution of workers across occupations and industries, rather than systematic differences in adoption within comparable jobs.

We find similar evidence when studying the relationship between partisanship and generative AI exposure. Democrats appear more exposed to generative AI in the raw data using the [Eloundou et al. \(2024\)](#) occupation level exposure index, but the association is largely absorbed by education and industry fixed effects. Put differently, partisan differences in exposure primarily reflect sorting into industries with higher predicted LLM exposure—which could be driven by different job and task requirements—rather than within industry exposure differences. In this sense, the observed “politics of AI” is better characterized as the political geography of education, occupations, and industries than as a direct ideological cleavage in technology adoption.

These patterns have implications for the political economy of technological change. If AI adop-

tion and AI intensive job opportunities are concentrated in specific industries and skill groups that correlate with party affiliation, the benefits of generative AI are likely to diffuse unevenly across political constituencies. Over time, uneven access to AI complementary jobs and organizational support for adoption could shape local labor market outcomes, alter perceived economic opportunity, and influence the politics of technology governance. At the same time, the weak partisan differences in perceived displacement risk suggest that differences in lived exposure and workplace integration may precede, rather than simply reflect, divergence in beliefs about AI.

This paper contributes to the political economy of technological change by documenting partisan stratification in workplace generative AI diffusion and by decomposing that stratification into composition and context. Prior work links exposure to economic shocks or automation to political outcomes, typically using regional measures of trade or robot exposure and election returns ([Autor et al., 2020](#); [Frey et al., 2018](#); [Anelli et al., 2019](#); [Petrova et al., 2025](#)). In contrast to these, which generally showed greater exposure of Republican-voting areas, we instead study a general purpose technology at the level where adoption occurs—the worker/workplace—combining self reported AI utilization with occupation based exposure measures and organizational AI environment indicators. The results show that raw partisan gaps in frequent use and task breadth are substantial and widening, yet they largely reflect education and industry occupational sorting, while within user differences in organizational support remain salient. This helps reconcile why groups can differ strongly in lived exposure and workplace integration even when partisan differences in perceived displacement risk are small, and it suggests a mechanism through which technology diffusion can shape future coalitions over AI governance ([Jacobs, 2024](#); [Anderson and Bishop, 2025](#)).

This paper also contributes to a quickly growing literature on measuring who uses generative AI and why, where the core challenge is that diffusion is fast and the relevant margin is behavior,

not exposure. Starting in 2023, [Gallup \(2026\)](#) began measuring AI use in the Gallup Workplace Panel—a longitudinal and nationally representative survey of U.S. adults and workers—finding that roughly 10% of respondents were using AI (broadly defined) at least a few times per week (21% if at all) in 2023 and 26% (46% if at all) by Q4:2025.³ [Bick et al. \(2025\)](#) provide a complementary measurement approach by fielding a high frequency survey designed to mirror the structure and labor market coverage of the Current Population Survey, finding a similar share of 23% reporting using generative AI for work at least once in the previous week (50% if at all). [Hartley et al. \(2024\)](#) have also been conducting surveys of AI use, finding a decline from 50% using generative AI as of August 2025 to 36% as of December 2025. A parallel strand of work measures AI adoption directly at the firm level. Using harmonized survey modules fielded across representative samples of firms in the United States, United Kingdom, Germany, and Australia, [Yotzov et al. \(2026\)](#) document that roughly 60-70% of firms report using at least one AI technology by late 2025—much higher than the 18% as of early 2026 reported by the Business Trends and Outlook Survey (BTOS) ([Bonney et al., 2024](#)). Crucially, however, the average executive reports only about 1.5 hours of personal AI use per week, consistent with micro-data on heterogeneous within-organization adoption ([Makridis, 2025](#)). Alternative approaches to measure AI at local ([Andreadis et al., 2025](#)) and firm ([Babina et al., 2024](#)) levels have relied upon job posting data.

2 Data and Measurement

This section describes the data underlying our analysis, explains the construction of the outcome variables, and presents the summary statistics.

³These numbers are also consistent with polling by Pew Research that found 58% of the population had heard of ChatGPT, but only 14% had tried it ([Vogels, 2023](#)).

The Gallup Workforce Panel. We analyze individual-level survey data from the Gallup Workforce Panel (GWP), a longitudinal dataset that provides rich measures of employee outcomes and characteristics. The GWP consists of a representative sample of U.S. adults who are surveyed about their work experiences on a recurring basis. In particular, we utilize the Gallup Panel over 2024 to 2026, covering 19,979 respondents and 47,778 person-waves.⁴ Gallup’s sampling methodology combines random digit dialing and address-based sampling to recruit panel participants, and survey weights are provided to ensure that the panel estimates are representative of the national workforce in terms of age, gender, region, and other demographics. In our analysis, we apply these weights so that our results generalize to the broader population of U.S. workers.

Outcome Variables. Our primary outcomes measure the use of AI for work. AI is defined broadly for respondents who are given the following definition: “Artificial intelligence is when a computer does or enhances a task that is normally performed by a person.” They are subsequently asked: “How often do you use artificial intelligence in your role,” which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. These responses are consolidated into three groups: often (daily and a few times a week), sometimes (a few times a month, a few times a year), and never (once a year, less often than once a year, and never), as well as into an ordinal AI utilization measure on a 0-10 scale: (Daily=10, Few Times Week=8, Few Times Month=6, Few times Year=4, Once a Year=2, Less than Once a Year=1, Never/Other=0). In some cases, we also predict political affiliation as a function of other covariates, including the [Eloundou et al. \(2024\)](#) measure of genAI exposure.

We gather four sets of responses on AI utilization: once in Q2:2024 when the AI module

⁴The GWP is available in prior years, but we focus on 2024 onward given the rapid pace of innovation with generative AI. While the GWP is quarterly, measurement of AI is only annual in 2023 before it turns quarterly in Q2:2024. We discuss these details further in the upcoming paragraphs.

was annual, and three times in 2025 when the module became quarterly in Q2:2025. Because of the rapid change in the adoption of Artificial Intelligence over this period, we provide results by sub-period as well as overall. Figure A.1 in the Online Appendix shows that roughly 70% of the respondents in the sample have at least two observations.

Generative AI Exposure. To proxy the technological opportunity to use AI in each job, we map the Eloundou et al. (2024) occupational large language model exposure scores—constructed by aligning O*NET tasks with LLM capabilities using expert and GPT-4 annotations—to respondents’ SOC-level codes. We aggregate the 6-digit scores to SOC-3 using the Occupational Employment and Wage Statistics (OEWS) employment records to maximize our Gallup Sample.⁵

Political Affiliation Political affiliation is measured through self-identification updated at least annually from the Gallup Panel with respondents identifying as Republican, Democrat, Independent.⁶ We restrict responses to those who report Republican or Democrat and allocate Independents to these two categories based on the following rule: if an individual reports Independent in time period t , we impute them as Republican or Democrat based on their most recent non-Independent affiliation in period $t - 1$. For those who only report being Independent, they are discarded from the sample. Doing so drops 25,225 responses (from 72,980).

Covariates. In addition to outcome variables, the data include detailed covariates that we use to control for potential differences across different political groups. Demographic variables in our dataset include age (in years), marital status, sex, race, full-time status, children in the household, educational attainment (indicators for highest degree earned, which we focus on some

⁵While we are able to map job titles into six-digit SOC codes for a third of the sample, we opt to maximize sample size and focus only on three-digit SOC codes, allowing us to keep roughly two-thirds of the sample for these occupational exposure exercises.

⁶The specific wording for the question is: ‘In politics, as of today, do you consider yourself a Republican, a Democrat, or an Independent?’ Other options include “Other party” and “Don’t know.”

college (including vocational) and college or more). Job-related variables include length of tenure with the current employer (measured in years), North American Industry Classification System (NAICS) codes, and Standard Occupational Classification (SOC) codes. Our baseline specification controls for age, sex, race, children in the household, education, and marital status to maximize sample size, but we also include state, SOC2 and NAICS2 fixed effects.

Summary Statistics. Table 1 documents the main descriptive statistics. We see that Democrats have 25% of responses that use AI frequently and 23% occasionally compared to 20% and 17% respectively for Republicans. In terms of demographics the Democrats are younger (average age of 42 vs 46 for Republicans), more gender balanced (46% male for Democrats vs 60% male for Republicans), less white, less likely to have children and much less likely to be married (49% for Democrats vs 66% for Republicans). One notable difference is in education with Democrats having 57% of respondees having college or more compared to 30% for Republicans.

3 Descriptive Results

Figure 2 shows a clear and persistent partisan gap in reported frequent AI utilization. In each survey wave, Democrats are more likely than Republicans to report using AI at least weekly, defined as answering “daily” or “a few times a week” to the question “How often do you use artificial intelligence in your role.” The gap is present in Q2:2024 and widens over time: 13.2% versus 9.6% in Q2:2024, 21.4% versus 15.4% in Q2:2025, 24.4% versus 19.8% in Q3:2025, 27.8% versus 22.6% in Q4:2025, and 30.1% versus 25.0% in Q1:2026. Figure A.3 reports a robustness check using a broader classification that separates “frequent users” (weekly or more) from “sometimes users” (monthly to yearly). The same qualitative pattern holds across both categories: Democrats exhibit

higher reported AI use in every wave, suggesting that the partisan gap reflects a general difference in adoption intensity rather than sensitivity to a particular cutoff in the frequency scale.⁷

Figure 3 plots average occupational exposure to generative AI using the task-based exposure index of Eloundou et al. (2024), mapped to respondents’ three-digit SOC codes. In each survey wave, Democrats are employed in occupations with slightly higher predicted exposure to AI than Republicans. The gap is present in Q2:2024 and remains stable through early 2026, with mean exposure levels of 0.465 versus 0.443 in Q2:2024, 0.472 versus 0.427 in Q2:2025, 0.462 versus 0.433 in Q3:2025, 0.460 versus 0.432 in Q4:2025, and 0.464 versus 0.433 in Q1:2026. Although the magnitude of the difference is modest in absolute terms, it is highly persistent and aligns with the utilization patterns in Figure 2. Thus, part of the partisan gap in AI adoption reflects systematic differences in occupational sorting, with Democrats more likely to work in roles whose task content is assessed as more amenable to large language model based automation or augmentation.

Next, Figure 4 reports perceived job displacement risk in response to the question “How likely is your job to be eliminated within 5 years due to technology or AI.” The figure plots the share of respondents who report that displacement is “somewhat likely” or “very likely.” Through 2025, partisan differences in perceived displacement risk were modest and statistically indistinguishable, with 14.1% of Democrats expressing concern about job loss compared to 13.5% of Republicans. However, by Q1:2026 the gap has widened considerably and become statistically significant, with 19.7% of Democrats reporting likely displacement compared to 15.5% of Republicans—a 4.2 percentage point difference ($p\text{-value} = 0.00$). This sharp increase in perceived displacement risk among

⁷Motivated by Jones (2026) who points out the rise of Independents, Figure A.2 in the Online Appendix deviates from our baseline approach of allocating Independents to Republican or Democrat categories based on past affiliation, instead treating them as their own group. Here, we find that Independents rank between Republicans and Democrats throughout the sample at 28.2% using AI at least multiple times per week as of Q1:2026, compared with 25.1% for Republicans and 30.2% for Democrats.

Democrats in 2026 is consistent with the differences observed in AI utilization and occupational exposure, and may reflect the fact that more frequent AI users—who skew Democratic—are more likely to perceive displacement risk. Importantly, [Makridis \(2026\)](#) also shows that workplace practices and managerial support can meaningfully offset these fears, suggesting that organizational context plays a central role in shaping how workers interpret AI adoption.

Finally, we construct an “AI Productivity Index,” consisting of ten indicators capturing whether respondents use AI for a range of work activities, including automating routine tasks, generating ideas, consolidating information, making predictions, and collaborating with others. [Figure 5](#) reports differences in the AI Productivity Index by political affiliation.⁸ Higher values therefore reflect both broader and more intensive use of AI across job functions. Democrats exhibit substantially higher scores on this index than Republicans, with mean values of 0.40 compared to 0.35. This gap is larger than what is observed for simple adoption measures and indicates that partisan differences extend beyond whether individuals use AI at all to how deeply AI is integrated into their work processes. Viewed in context of the utilization and occupational exposure results, this pattern suggests that Democrats are not only more likely to be in AI-exposed roles, but also more likely to deploy AI across a wider set of productivity-relevant tasks within those roles.

In summary, the figures point to a consistent partisan gradient in workplace AI engagement. Democrats are more likely to use AI frequently, work in occupations with higher predicted AI exposure, and integrate AI more deeply across a range of productivity relevant tasks. These differences are persistent over time and robust to alternative usage thresholds, suggesting that they

⁸Figure [A.4](#) in the Online Appendix examines the partisan differences by category. Conditional on using AI at least somewhat during the year, Democrats are more likely to report using AI for higher level cognitive and information processing tasks, including generating ideas (50.5% vs. 42.5%) and consolidating information (51.7% vs. 44.3%). Republicans report relatively higher use for operational and support activities such as learning new things (46.7% vs. 42.6%), interacting with customers (22.0% vs. 14.4%), and making predictions (14.5% vs. 8.6%). Thus, there are partisan differences not only in adoption intensity, but also in the functional composition of AI use.

reflect both occupational sorting and differential intensity of adoption. In contrast, while partisan differences in perceived job displacement risk were small and imprecisely estimated through 2025, but a notable gap of 4 percentage points has emerged by Q1:2026, consistent with the broader divergence in AI utilization and exposure. In the next section we investigate to what extent this is due to educational, occupational and industrial differences between these groups.

4 Regression Results

To investigate the drivers of the higher Democrat exposure to AI, we also run a series of linear models for both outcomes on pooled cross-sections over 2024-2025:

$$y_{i,t} = \alpha + \beta \text{Dem}_i + X_i' \gamma + \lambda_t + \delta_s + \theta_o + \kappa_n + \varepsilon_{i,t}, \quad (1)$$

where $y_{i,t}$ is either *Use AI Often* (LPM) or *AI Use Intensity*, Dem_i is an indicator for Democratic affiliation, X_i includes demographics (age, sex, race, marital status, children, and educational attainment), λ_t are year-quarter fixed effects, and $(\delta_s, \theta_o, \kappa_n)$ are state, SOC2, and NAICS2 fixed effects, respectively. We cluster standard errors at the person-level.

To proxy technological opportunity at the job level, we merge the occupation-level exposure measure from [Eloundou et al. \(2024\)](#), who map O*NET tasks to LLM capabilities using a rubric implemented by human annotators and aggregate to SOC-3, with our individual-level data:

$$\text{Exposure}_o = \alpha + \rho \text{Dem}_{i,t} + X_i' \gamma + \lambda_t + \delta_s + \kappa_n + \varepsilon_{i,t}. \quad (2)$$

This specification asks whether (i) partisan affiliation explains raw exposure and (ii) any residual

exposure varies with partisanship. We report results using the standardized continuous index.

Table 2 reports linear probability models for whether respondents use AI frequently, defined as daily or a few times per week. In the baseline specification with only time fixed effects, Democrats are 4.9 pp more likely than Republicans to be frequent AI users. This gap remains statistically significant after controlling for age and basic demographics, but is sharply attenuated once educational attainment (namely college attainment) is added, falling to 0.6 pp and becoming statistically insignificant. Introducing occupational or industry fixed effects also reduces the estimate, and in the fully saturated model with both SOC2 and NAICS2 fixed effects the coefficient becomes negative and statistically significant. This pattern indicates that the raw partisan gap in frequent AI use is largely explained by compositional differences in education and job characteristics, rather than within occupation differences in adoption behavior, and that Democrats are actually 2.2 pp less likely to use AI after adjusting for demographics and job-level differences.

Table A.1 provides a robustness check using the continuous AI Use Intensity index on a 0 to 10 scale. The results mirror those for frequent use. Democrats report substantially higher intensity in the baseline, with a difference of 0.70 points, and the gap remains positive but smaller after adding age and basic demographics. However, once education and job fixed effects are included, the coefficient declines sharply and becomes small or negative in the fully specified model. A main reason for the more attenuated fully saturated specification stems from the partisan differences arising more from frequent AI users, rather than the sometimes users (Figure A.3).

Next, Table 3 estimates the association between party affiliation and occupational exposure to generative AI, measured as the standardized Eloundou et al. (2024) LLM β exposure index mapped to respondents' occupations. In the parsimonious specifications, Democrats work in occupations with significantly higher predicted LLM exposure, with coefficients of 0.183 and 0.202 standard

deviations in columns (1) and (2). However, once NAICS2 industry fixed effects and education are added, the Democrat coefficient falls sharply to 0.036 and then 0.032 and is no longer statistically distinguishable from zero in columns (3) and (4). This attenuation indicates that the partisan differences in occupational exposure are largely explained by industry composition and education rather than systematic differences in exposure within industries. These patterns are consistent with sorting by demographic and job characteristics: higher education is associated with higher exposure, while males and respondents with children tend to be in lower exposure roles, and tenure is positively associated with exposure in the specification where it is included.

Finally, Table 4 compares workplace AI related indicators between Democrats and Republicans. Democrats exhibit higher engagement and more favorable workplace conditions for AI adoption across nearly all dimensions. Democrats are more likely to use AI frequently (25.1% vs. 20.2%) and report higher scores on the AI Productivity Index (0.197 vs. 0.142), indicating broader and more intensive task level integration. Democrats also report more supportive organizational environments. They are more likely to work in firms that have integrated AI tools (44.5% vs. 34.9%) and communicated a clear AI strategy (25.8% vs. 21.9%). They report higher preparedness and comfort using AI, and greater access to optional training and informal advice. In contrast, required AI training does not differ across parties. Overall, Democrats operate in organizational contexts that are more conducive to sustained and advanced AI adoption, but the fact that these correlations flip after controlling for demographic and job-level factors in Table 2 suggests that the raw correlation is related to compositional differences.

5 Conclusion

This paper provides the first evidence on how AI in the workplace is diffusing across the U.S. labor market along political lines. Using nationally representative Gallup Workforce Panel data linked to occupation-level measures of AI exposure from Eloundou et al. (2024) through Q1:2026, we document a persistent partisan gradient in AI adoption and task-level integration. In the raw data, Democrats are more likely to use AI frequently, operate in more AI-exposed occupations, and deploy AI across a wider range of productivity-relevant tasks. However, these differences are largely explained by educational attainment and sorting into industries and occupations where generative AI is more technologically relevant. Once these compositional factors are accounted for, within-job partisan differences in adoption largely disappear or reverse.

These findings suggest that the “politics of AI” is not primarily driven by ideological resistance or enthusiasm for the technology, but rather by structural differences in where people work and what skills they possess. Political affiliation appears to proxy for deeper patterns of human capital accumulation, occupational clustering, and organizational environment. In this sense, the distribution of AI benefits across political groups may reflect longstanding inequalities in access to high-skill, technology-complementary jobs more than ideological differences in individual preferences toward AI itself. From a political economy perspective, this has important implications. If AI complements high-education and professional occupations, and these remain unevenly distributed across political constituencies, then the economic gains from AI may reinforce existing partisan cleavages in income growth, skill accumulation, and workplace opportunity. Over time, these material differences in lived exposure and organizational support may shape how different

groups evaluate AI regulation, labor policy, and innovation more broadly, feeding into broader political coalitions—particularly as the emerging partisan gap in perceived job displacement risk suggests that belief differences are beginning to widen alongside behavioral ones.

Our analysis opens several directions for future work. First, linking individual AI adoption to local labor market outcomes and voting behavior could clarify whether workplace exposure to generative AI translates into observable political responses, as has been shown for trade and industrial automation. Second, as the Gallup Workforce Panel continues, we will gather even richer longitudinal data that enables studying dynamic career effects, such as whether AI users experience faster wage growth, occupational mobility, or changes in job stability, and whether these trajectories differ by political affiliation. Third, future studies could examine firm-level heterogeneity in AI strategy, governance, and training investments to better understand how organizational choices mediate the distribution of AI benefits. Fourth, extending this analysis beyond the U.S. would help assess whether similar partisan or ideological gradients emerge in countries with different political systems, labor market institutions, and technology adoption environments. Fifth, given the rise of Independents ([Jones, 2026](#)), it will be interesting to explore their political experiences and entry in and out of Republican and Democrat affiliations over time.




Figure 1: Media Coverage of AI Economics Has Surged

Notes.—Sources: AI Economic Index, measuring the coverage of AI and the Economy in around 2000 daily newspapers in the US, normalized to 100 from 2015 to 2021 inclusive. Coverage defined as articles that included (AI, artificial intelligence, genAI, machine learning, computer vision) AND (economic, economy). Based on the methodology of Baker et al. (2016) and provided on https://www.policyuncertainty.com/aieu_daily.html.




Figure 2: Partisan Differences in AI Utilization

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the share of respondents who answer “daily” or “a few times a week” to the question “How often do you use artificial intelligence in your role,” which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. Before asking this question, respondents are given a definition of AI: “Artificial intelligence is when a computer does or enhances a task that is normally performed by a person.” Then, we distinguish between Republicans and Democrats. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.




Figure 3: Occupational Exposure by Partisanship

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the three-digit occupational exposure from the [Eloundou et al. \(2024\)](#) index by Republican and Democrat affiliation. The β exposure is an occupation-level index of LLM exposure constructed by classifying O*NET tasks into no exposure, direct exposure (E1) and LLM+ exposure (E2) based on whether access to GPT-4 alone or with complementary software would reduce the time to complete the task by at least 50% at constant quality, using both human annotators and GPT-4 as a classifier. The β index aggregates these labels to occupations as $\beta = E1 + 0.5 \times E2$, so that directly exposed tasks are counted twice as heavily as tasks that require complementary tools. We map these occupation-level scores to respondents' SOC codes. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.




Figure 4: Perceptions of Job Displacement from AI by Partisanship

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the share of respondents who report “somewhat likely” or “very likely” in response to the question: “How likely is your job to be eliminated within 5 years due to technology or AI.” In Q1:2026, Democrats exhibit a 4.2 percentage point higher probability of believing their job will be displaced (p -value = 0.00); the share of respondents reporting somewhat or very likely is 12%. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.




Figure 5: AI Productivity Index, by Partisanship

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the AI Productivity Index by party affiliation, which is defined as the sum (0-10) of ten indicators for whether the respondent reports using AI for specific work activities—automating basic tasks, identifying problems, learning new things, setting up/operating/monitoring devices, collaborating with others, interacting with customers, generating ideas, consolidating information, making predictions, or other—counting an activity only when the corresponding item equals 1 (zeros, i.e. non-AI users, are also retained in the sample). The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Table 1: Summary Statistics on the Gallup Workforce Panel

	Republican	Democrat
	(1)	(2)
	mean	mean
Frequent AI Use	0.20	0.25
Occasional AI Use	0.17	0.23
Never Use AI	0.62	0.52
AI Utilization	2.8	3.5
AI Will Displace Job	0.14	0.17
Hours Worked/Week	40.8	39.6
Age	46.1	42.3
Male	0.60	0.46
White	0.78	0.58
Black	0.03	0.17
Full-time	0.87	0.87
Part-time	0.13	0.13
Has Children	0.40	0.34
Married	0.66	0.49
Less than HS	0.41	0.20
Some College	0.29	0.23
College or More	0.30	0.57
Always Remote	0.11	0.16
Hybrid Remote	0.18	0.32
Never Remote	0.71	0.52
Share of Sample	0.41	0.59
Observations	17470	30308

Notes.—Sources: Gallup Panel, 2024-2026. The table reports summary statistics on the key variables by political affiliation. AI Use - Often is based on using AI daily or a few times a week; AI Use - Sometimes is based on using AI a few times a month or a few times a year; AI Use - Never is based on using AI once a year, less than once a year, or never. The AI frequency count is defined as follows: Daily=10, Few Times Week=8, Few Times Month=6, Few times Year=4, Once a Year=2, Less than Once a Year=1, Never/Other=0. AI Will Displace Job is an indicator for believing that AI is “very likely” or “somewhat likely” to replace their job in the next five years. Demographics include: age, race (White, Black), full-time status, having children, and being married. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Table 2: Political Affiliation as a Predictor of Frequent AI Use

Dep. var. =	Use AI Often (1/0)					
	(1)	(2)	(3)	(4)	(5)	(6)
Democrat	.049*** [.007]	.042*** [.008]	.006 [.007]	.013* [.008]	.023*** [.007]	-.022*** [.008]
Age		-.003*** [.000]				-.002*** [.000]
Some College			.055*** [.010]			.015 [.011]
College or More			.173*** [.009]			.091*** [.012]
Male		.025*** [.007]				.029*** [.008]
White		-.038*** [.010]				-.045*** [.009]
Black		.011 [.014]				.040*** [.015]
Has Children		.012 [.008]				.021*** [.008]
Is Married		.050*** [.008]				.013* [.008]
R-squared	.02	.03	.05	.08	.07	.12
Sample Size	47778	47296	47768	39994	45454	37688
SOC2 FE	No	No	No	Yes	No	Yes
NAICS2 FE	No	No	No	No	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Sources: Gallup Panel, 2024-2026. The table reports the coefficients associated with regressions of an indicator for whether the respondent uses AI daily or frequently during the week (zero otherwise) on an indicator for affiliating as a Democrat, a vector of individual covariates and layers of fixed effects. Demographics include: indicators for some college (two-year, vocational, or some college) and college or more, age, race (White, Black), having children, and being married. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Table 3: Predicting Political Affiliation Based on Generative AI Exposure

Dep. var. =	LLM β Exposure (z-score)			
	(1)	(2)	(3)	(4)
Democrat	.183*** [.032]	.202*** [.033]	.036 [.030]	.032 [.030]
Age		-.001 [.001]	.002 [.001]	-.000 [.001]
Some College			.209*** [.045]	.212*** [.046]
College or More			.171*** [.048]	.178*** [.049]
Male		-.025 [.030]	-.151*** [.026]	-.156*** [.026]
White		-.028 [.037]	-.043 [.032]	-.050 [.032]
Black		-.089* [.050]	-.076* [.042]	-.079* [.042]
Has Children		-.149*** [.033]	-.106*** [.028]	-.102*** [.028]
Is Married		.164*** [.032]	.092*** [.027]	.086*** [.027]
Tenure				.007*** [.002]
R-squared	.01	.02	.28	.28
Sample Size	26545	26267	24991	24333
Demographics	No	Yes	Yes	Yes
NAICS2 FE	No	No	Yes	Yes

Notes.—Sources: Gallup Panel, 2024-2026, and Eloundou et al. (2024) The table reports the coefficients associated with regressions of the Eloundou et al. (2024) β exposure on an indicator for whether the respondent is Democrat, a vector of individual covariates and layers of fixed effects and job tenure. The β exposure is an occupation-level index of LLM exposure constructed by classifying O*NET tasks into no exposure, direct exposure (E1) and LLM+ exposure (E2) based on whether access to GPT-4 alone or with complementary software would reduce the time to complete the task by at least 50% at constant quality, using both human annotators and GPT-4 as a classifier. The β index aggregates these labels to occupations as $\beta = E1 + 0.5 \times E2$, so that directly exposed tasks are counted twice as heavily as tasks that require complementary tools, and we map these occupation-level scores to respondents' SOC codes. Demographics include: indicators for some college (two-year, vocational, or some college) and college or more, age, race (White, Black), having children, being married, and job tenure. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Table 4: AI Workplace Indicators by Party

	Democrat	Republican	Diff (D–R)	p-value	N Dem	N Rep
Uses AI often	0.251	0.202	0.049	0.000	30308	17470
Uses AI sometimes	0.225	0.174	0.051	0.000	30308	17470
AI Productivity Index	0.197	0.142	0.054	0.000	30308	17470
Believes job likely displaced by tech/AI	0.171	0.143	0.028	0.000	14188	8138
Firm has integrated AI tools	0.445	0.349	0.096	0.000	14400	8271
Firm communicated clear AI strategy	0.258	0.219	0.039	0.000	28598	16468
Prepared to work with AI/advanced tech	0.385	0.316	0.069	0.000	3173	1643
Comfort using AI in role	0.188	0.148	0.040	0.000	3479	1982
Required AI training	0.010	0.013	-0.002	0.026	30308	17470
Optional AI training opportunities	0.024	0.016	0.008	0.000	30308	17470
Informal AI advice/tips	0.025	0.018	0.008	0.000	30308	17470

Notes.—Sources: Gallup Panel, 2024–2026. The table reports unweighted group means for Democrats and Republicans, the mean difference (Dem–Rep), the two-sided p-value from a Welch two-sample t-test (unequal variances), and group-specific non-missing sample sizes for each outcome (N Dem and N Rep). Variable definitions: *Uses AI often* is an indicator equal to 1 if the respondent reports using AI daily or a few times a week (Q619A categories 1–2) and 0 otherwise. *Uses AI sometimes* is an indicator equal to 1 if the respondent reports using AI a few times a month or a few times a year (Q619A categories 3–4) and 0 otherwise. *AI Productivity Index* is defined as the sum (0–10) of ten indicators for whether the respondent reports using AI for specific work activities: automating basic tasks, identifying problems, learning new things, setting up/operating/monitoring devices, collaborating with others, interacting with customers, generating ideas, consolidating information, making predictions, or other—counting an activity only when the corresponding item equals 1 (zeros, i.e. non-AI users, are also retained in the sample). *Believes job likely displaced by tech/AI* is an indicator equal to 1 if the respondent believes AI/advanced technologies are somewhat likely or very likely to replace their job in the next five years (ai_likelydisplace categories 3–4) and 0 otherwise. *Organization has integrated AI tools* is an indicator equal to 1 if the respondent reports their organization has integrated AI technology or tools (harmonized from Q653/Q653A/Q653B depending on quarter) and 0 otherwise; “Don’t know” responses (where available) are coded as 0. *Organization communicated clear AI strategy/plan* is an indicator equal to 1 if the respondent reports that their organization has communicated a clear plan or strategy for integrating AI and 0 otherwise (Q654A). *Prepared to work with AI/advanced technologies* is an indicator created when equal to 3 or 4 (somewhat or very prepared) from a 1–4 scale capturing how prepared the respondent feels to work with AI, robotics, or other advanced technologies, where higher values indicate greater preparedness (Q618). *Comfort using AI in role* is a similar indicator created when equal to 4 or 5 (somewhat or very comfortable) on a 1–5 scale capturing the respondent’s comfort with using AI in their role, where higher values indicate greater comfort (Q657). Training measures are indicators based on a multi-response item about AI training at work (Q659): *Required AI training* equals 1 if AI training is required and 0 otherwise; *Optional AI training opportunities* equals 1 if optional AI training is offered and 0 otherwise; *Informal AI advice/tips* equals 1 if informal advice/tips are available and 0 otherwise. All variables are coded missing when the underlying survey item is missing. The sample is restricted to respondents ages 25–65 and either full or part time.

A Online Appendix




Figure A.1: Descriptive Evidence on Longitudinal Variation

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the share of respondents across the potential number of waves in the data.

Table A.1: Political Affiliation as a Predictor of AI Utilization

Dep. var. =	AI Use Intensity (0-10)					
	(1)	(2)	(3)	(4)	(5)	(6)
Democrat	.699*** [.073]	.643*** [.076]	.133* [.071]	.229*** [.073]	.413*** [.072]	-.099 [.074]
Age		-.041*** [.003]				-.022*** [.003]
Some College			.667*** [.099]			.175* [.101]
College or More			2.244*** [.092]			1.121*** [.107]
Male		.313*** [.070]				.405*** [.071]
White		-.375*** [.094]				-.503*** [.084]
Black		-.098 [.134]				.165 [.130]
Has Children		.234*** [.076]				.380*** [.072]
Is Married		.586*** [.074]				.127* [.070]
R-squared	.03	.05	.09	.15	.10	.20
Sample Size	47778	47296	47768	39994	45454	37688
SOC2 FE	No	No	No	Yes	No	Yes
NAICS2 FE	No	No	No	No	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Sources: Gallup Panel, 2024-2026. The table reports the coefficients associated with regressions of a frequency count for AI utilization on an indicator for affiliating as a Democrat, a vector of individual covariates and layers of fixed effects. The frequency count is defined as follows: Daily=10, Few Times Week=8, Few Times Month=6, Few times Year=4, Once a Year=2, Less than Once a Year=1, Never/Other=0. Demographics include: indicators for some college (two-year, vocational, or some college) and college or more, age, race (White, Black), having children, and being married. The sample is restricted to full and part time workers, and the observations are weighted.




Figure A.2: Partisan Differences in AI Utilization (With Independents)

Notes.—Sources: Gallup Panel, 2024-2026. The figure plots the share of respondents who answer “daily” or “a few times a week” to the question “How often do you use artificial intelligence in your role,” which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. Before asking this question, respondents are given a definition of AI: “Artificial intelligence is when a computer does or enhances a task that is normally performed by a person.” We allow for the full sample of Republicans, Democrats, and Independents. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.




Figure A.3: Partisan Differences in AI Utilization, 2024-25

Notes.—Sources: Gallup Panel, 2024-2025. The figure plots the share of respondents who answer “daily” or “a few times a week” in the “active category” and separately “a few times a month” to “once a year” in the “sometimes” category to the question “How often do you use artificial intelligence in your role,” which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. Before asking this question, respondents are given a definition of AI: “Artificial intelligence is when a computer does or enhances a task that is normally performed by a person.” Then, we distinguish between Republicans and Democrats. The sample is restricted to full and part time workers, and the observations are weighted.




Figure A.4: Composition of AI Activities Among Democrats and Republicans

Notes.—Sources: Gallup Panel, 2026. The figure plots the share of respondents who report different AI use-cases and activities, conditional on using AI at least somewhat during the year. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

References

- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):293–340.
- Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation, and work. *Journal of Economic Perspectives*, 32(2):197–236.
- Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.
- Anderson, M. and Bishop, W. (2025). Republicans, democrats now equally concerned about ai in daily life, but views on regulation differ. Pew Research Center, Short Reads.
- 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.
- Anelli, M., Colantone, I., and Stanig, P. (2019). We were the robots: Automation and voting behavior in western europe. IZA Discussion Paper 12485, IZA Institute of Labor Economics.
- Autor, D., Dorn, D., Hanson, G. H., and Majlesi, K. (2020). Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–3183.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

- Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151(103745).
- Baker, S., Bloom, N., and Davis, S. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.
- Bick, A., Blandin, A., and Deming, D. J. (2025). The rapid adoption of generative ai. Working Paper 2024-027, Federal Reserve Bank of St. Louis.
- Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L., 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. Ces working paper no. ces-24-16 / nber working paper no. 32319, U.S. Census Bureau, Center for Economic Studies.
- Brynjolfsson, E., Chandar, B., and Chen, R. (2025a). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab.
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2025b). Generative ai at work. *Quarterly Journal of Economics*, 140(2):889–942.
- 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.
- 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.

- Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.
- Frey, C. B., Berger, T., and Chen, C. (2018). Political machinery: Did robots swing the 2016 u.s. presidential election? *Oxford Review of Economic Policy*, 34(3):418–442.
- Gallup (2026). What percentage of employees use ai at work? *The Gallup Organization*.
- Hall, A. B., Westwood, S. J., and Grimmer, J. (2025). Measuring perceived slant of large language models through user evaluations. Stanford GSB Working Paper.
- Hartley, J., Jolevski, F., Melo, V., and Moore, B. (2024). The labor market effects of generative artificial intelligence. *Available at SSRN*.
- Humlum, A. and Vestergaard, E. (2025). Large language models, small labor market effects. Working Paper 33777, National Bureau of Economic Research.
- Jacobs, J. (2024). The artificial intelligence shock and socio-political polarization. *Technological Forecasting and Social Change*, 199:123006.
- Johnston, A. C. and Makridis, C. A. (2025). The labor market effects of generative ai: A difference-in-differences analysis of ai exposure. Working Paper, available at SSRN (No. 5375017).
- Jones, J. (2026). New high of 45% in u.s. identify as political independents. *Gallup*.
- Makridis, C. A. (2025). Organizational transmission of ai: Managerial influence on generative ai adoption. *SSRN working paper*.

- Makridis, C. A. (2026). The sum of all (workplace) fears: How managers mediate the fear of ai job displacement. *SSRN working paper*.
- Mills, M. A. and St. Clair, P. (2025). The strange new politics of science. *Issues in Science and Technology*, 41(3). Spring.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.
- Petrova, M., Schubert, G., Taska, B., and Yildirim, P. (2025). Robotization and the political response of politicians. In *The Political Economy of Artificial Intelligence*, chapter 10. University of Chicago Press.
- Vogels, E. A. (2023). A majority of americans have heard of chatgpt, but few have tried it themselves. *Pew Research Center*.
- Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K., Jalca, A., Meyer, B., Mizen, P., Navarrete, M., Smietanka, P., Thwaites, G., and Wang, B. Z. (2026). Firm data on AI. NBER working paper 34836.