

BLS WORKING PAPERS



U.S. Department of Labor
U.S. Bureau of Labor Statistics
Office of Productivity and Technology

AI Exposure and Adoption among U.S. Firms

Matthew Dey

U.S. Bureau of Labor Statistics

Emin Dinlersoz

U.S. Census Bureau

Lucia Foster

U.S. Census Bureau

Cheryl Grim

U.S. Census Bureau

John Haltiwanger

University of Maryland

Peter B. Meyer

U.S. Bureau of Labor Statistics

Sabrina Wulff Pabilonia

U.S. Bureau of Labor Statistics

Aditya Pande

U.S. Census Bureau

Martha Stinson

U.S. Census Bureau

Sean Wang

U.S. Census Bureau

Zoltan Wolf

BLN24

This paper is also posted under the same title as CES Working Paper 26-61

Working Paper 591

September 2026

AI Exposure and Adoption among U.S. Firms

Matthew Dey^a Emin Dinlersoz^b Lucia Foster^c
Cheryl Grim^b John Haltiwanger^d Peter B. Meyer^a
Sabrina Wulff Pabilonia^a Aditya Pande^{b,d} Martha Stinson^b
Sean Wang^b Zoltan Wolf^e

September 25, 2026

Abstract

Measures of exposure to artificial intelligence (AI) are widely used to study where AI is likely to affect firm and worker outcomes, yet limited evidence exists on how closely exposure relates to realized AI adoption at the firm level. We examine this relationship by linking 14 firm-level exposure measures, constructed from occupational exposure estimates in the literature and occupational employment shares from the Bureau of Labor Statistics Occupational Employment and Wage Statistics program, to direct measures of firm AI adoption from the Census Bureau’s Business Trends and Outlook Survey. Exposure is positively and statistically significantly associated with adoption, but explains only a modest share of its variation. The strength of the relationship varies considerably across standardized exposure measures: a one-standard-deviation increase in firm-level exposure is associated with a 4–11 percentage point higher firm adoption probability, falling to 1–8 percentage points after controlling for year and sub-sector fixed effects. The exposure–adoption relationship is also heterogeneous across firm-size classes and sectors. Exposure is thus an informative but incomplete signal of adoption. More broadly, the results have implications for the interpretation of exposure-based measures in studies of the economic effects of AI.

JEL Classification: O33, D22, J23

Keywords: Artificial Intelligence, AI adoption, AI exposure, firm-level data, occupational mix

^a U.S. Bureau of Labor Statistics ^b U.S. Census Bureau ^c U.S. Census Bureau (retired) ^d University of Maryland ^e BLN24
Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the U.S. Bureau of Labor Statistics. The authors used AI tools to edit parts of this paper for clarity; the authors take full responsibility for the final text. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure-avoidance protection (Project 7526914: CBDRB-FY26-CES007-02). Bureau of Labor Statistics staff reviewed the results to ensure respondent confidentiality. We thank Barbara Downs, Julie Hatch, Anne Polivka, and Lee Tucker for their helpful comments. The project was started when Lucia Foster was a Census Bureau employee. At the time this paper was written, John Haltiwanger was also a part-time employee of the U.S. Census Bureau. Emails: {dey.matthew,meyer.peter,pabilonia.sabrina}@bls.gov, {emin.m.dinlersoz,cheryl.ann.grim,martha.stinson,sean.y.wang,zoltan.wolf}@census.gov, lucia.foster.economics@gmail.com, {halt,apande}@umd.edu.

1 Introduction

A growing literature focuses on the adoption and economic consequences of AI. A central challenge in this work is measuring where AI is likely to matter for workers and firms. Many researchers have addressed this challenge by constructing measures of AI exposure based on workers’ tasks and occupations.¹ These measures provide a systematic way to characterize the technological relevance of AI across occupations and industries and have been used to study a wide range of labor-market and firm outcomes.² At the same time, exposure and adoption (or use) capture different margins of AI-driven technological change.³ Exposure measures the extent to which AI is applicable to the tasks performed by workers or occupations, whereas adoption reflects whether firms convert those technological opportunities into actual AI use. Adoption is a decision variable that depends not only on exposure, but also on other factors, such as barriers to adoption (e.g., implementation costs), organizational capabilities, complementary inputs, and other firm-specific factors. For example, a firm may adopt AI because it can be applied to its software, equipment, and capital-intensive processes, which are not necessarily captured by occupational exposure measures.

How closely exposure and adoption align is an empirical question with implications for both measurement and interpretation. If the two are tightly linked, exposure measures are informative not only about where AI is technologically applicable but also about where diffusion is actually occurring. However, if the link is weak, or varies across firms and industries, exposure and adoption capture separate margins of technological change, and an association between exposure and firm or worker outcomes cannot be read as the consequence of realized AI use without further evidence on adoption. Firm-level evidence on this question

¹See, e.g., Eloundou et al. (2024), Felten et al. (2018), Felten et al. (2021), Massenkoff and McCrory (2026).

²See, e.g., Acemoglu et al. (2022), Brynjolfsson et al. (2026), Eisfeldt et al. (2026), Frank et al. (2025), Hampole et al. (2025), Tucker (2026).

³Throughout the paper, we use AI adoption and AI use interchangeably. Empirically, adoption is measured by a binary indicator of firms’ reported use of AI in the survey data.

has been scarce, largely because comprehensive measures of firm AI use have not been available alongside occupational exposure measures.

This paper provides firm-level evidence on the relationship between AI exposure and adoption by linking measures of firms’ AI exposure to their reported AI use and examining how the two vary across U.S. firms. We construct firm-level exposure measures by combining several occupation-level AI exposure measures from the literature with establishment-level occupational employment shares from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS) survey. Measures of firm-level AI adoption come from the Census Bureau’s Business Trends and Outlook Survey (BTOS). Linking these data to the Census Bureau’s Longitudinal Business Database (LBD) provides firm characteristics used in the analysis, such as industry, size, age, and location.

We restrict the analysis to single-unit firms because their occupational composition can be mapped directly from establishment-level OEWS information to firm-level AI adoption. For multi-unit firms, the OEWS may cover only a subset of the firm’s establishments, so the observed occupational mix does not necessarily represent the workforce exposure of the entire firm. While this restriction ensures the comparability of the exposure and adoption measures at the firm level, it excludes many large firms that also tend to have relatively high rates of AI adoption (Bonney et al., 2024, 2026). The results should therefore be interpreted as characterizing the exposure–adoption relationship among single-unit firms rather than being representative of all firms.

We use the resulting dataset to characterize the relationship between AI exposure and adoption in a regression framework, without assigning a causal interpretation to the estimated associations. We examine this relationship for 14 exposure measures and document the findings in more detail for two exposure measures commonly used in the literature: the GPT-4 based beta exposure measure in Eloundou et al. (2024) and the observed exposure measure in Massenkoff and McCrory (2026). We explore heterogeneity in this relationship by sector and firm size, and conduct some robustness checks.

Several findings emerge from the analysis. First, exposure is systematically associated with realized firm adoption, but the strength of this relationship varies considerably across exposure measures. Moreover, exposure explains only a small share of the cross-firm variation in adoption. All 14 exposure measures yield positive and statistically significant unconditional associations with adoption. When the exposure measures are standardized, a one-standard-deviation increase in exposure is associated with roughly a 4–11 percentage point higher probability of adoption. However, these relationships attenuate substantially after conditioning on observables. Including year and three-digit industry fixed effects reduces the corresponding range to approximately 1–8 percentage points. For instance, for both the GPT-4 beta exposure measure of Eloundou et al. (2024) and the observed exposure measure of Massenkoff and McCrory (2026), the exposure coefficient falls by roughly one-half when detailed industry controls are introduced. Exposure alone accounts for only about 1–6 percent of the variation in adoption across firms, and its explanatory power within detailed industries is much smaller. Thus, although exposure has some predictive content for where AI adoption is occurring, substantial heterogeneity in adoption remains among firms with similar levels of exposure.

Second, the exposure–adoption relationship varies across sectors and firm types. The Information sector stands out as having both the highest-adoption rate and the sector where occupational exposure is most strongly associated with adoption, but some other adoption-heavy sectors, such as Education Services, do not have a strong relationship. Several other sectors exhibit much weaker relationships, including some indistinguishable from zero. Adoption rates predicted by occupational exposure alone compress the observed differences across sectors, tending to under-predict adoption in high-adoption sectors such as Information and Professional Services and to over-predict it in several low-adoption sectors. The point estimates for the exposure-adoption relationship also tend to be higher among larger firms—though our estimates in the largest size class are less precise. These patterns indicate that industry and firm characteristics contain important information about realized adoption

beyond what is contained in occupational exposure. Furthermore, the main findings are robust to restricting the sample to higher-quality matches, using firm weights rather than employment weights, and measuring adoption with different versions of the BTOS AI use question.

The findings overall suggest that exposure and adoption capture related but distinct dimensions of AI diffusion, with exposure providing an informative but incomplete signal of adoption. As a result, an exposure–outcome relationship cannot generally be interpreted as an adoption–outcome relationship without further assumptions or evidence on adoption. Instead, the exposure–outcome relationships are more naturally interpreted as likely reflecting the effects of technological opportunity or susceptibility, which may operate in part through realized adoption as well as through other channels, such as firms’ anticipation of, uncertainty about, and adjustments to AI diffusion. This interpretation may be especially relevant given the current early stage of AI diffusion across firms, where firm-level adoption remains low and factors other than adoption may be more relevant for some outcomes. Exploring the separate contributions of exposure and adoption to outcomes, and the exact mechanisms by which they contribute, is therefore an important and challenging future research avenue for the interpretation of exposure–outcome relationships. For settings where adoption data are available along with exposure measures, exploiting the variation across firms or industries with different adoption and exposure rates can help distinguish effects associated directly with AI use from those associated mainly with exposure.

Our findings have implications for a large and growing literature that aggregates occupation- and task-level AI exposure measures to the establishment, firm, or industry level in order to study how differences in exposure are associated with firm or worker outcomes (see, e.g., Acemoglu et al. (2022), Eisefeldt et al. (2026), Eloundou et al. (2024), Hampole et al. (2025), and Labaschin et al. (2025)). Together, these studies demonstrate how occupational and task-based measures can be used to characterize firms’ exposure to AI and relate that exposure to economic outcomes. They generally do not, however, assess whether and how

closely such exposure measures correspond to firms’ realized AI use using direct measures of adoption, in part due to the lack of comprehensive and representative adoption measures at the firm-level. Our analysis enables this comparison by linking occupation-based measures of firm AI exposure to direct measures of firms’ reported AI use in a large representative survey (BTOS).

A small recent literature examines the exposure-adoption relationship directly, at levels of aggregation ranging from sectors to occupations to firms. Tucker (2026) relates adoption measures from public-use BTOS data to exposure at the sector and sub-sector levels, but not at the firm level. A one standard-deviation increase in sub-sector AI exposure based on the Eloundou et al. (2024) GPT-4 beta measure is associated with a 6.7-percentage-point increase in adoption, with 47% of the observed variation in adoption explained.⁴ Industry aggregation strengthens the exposure–adoption relationship by averaging out firm-level idiosyncratic factors, but it can also obscure substantial adoption heterogeneity among exposed firms. Firm-level analysis matters because adoption is a firm decision and informs the interpretation of exposure–outcome relationships. For instance, an association between industry exposure and employment or productivity need not imply that the firms driving the association are the ones that actually adopted AI. The association may arise due to an adoption channel, an exposure channel that operates independently of adoption, or both.

Bick et al. (2026) use occupation-level generative AI adoption data collected from workers as part of their Real-Time Population Survey to compare with six generative AI exposure measures including three from Eloundou et al. (2024). Four results that are most salient for our work are as follows. First, there are positive correlations between exposure measures and adoption. Second, the correlations are well below one, indicating gaps between exposure and adoption. Third, exposure measures account for less than half of the variation in adoption

⁴This result comes from a bivariate regression of exposure on adoption. Across different exposure measures, the estimated sub-sector (3-digit NAICS) level exposure–adoption coefficient is in the range of 4.1–7.7 percentage points, with 17–58% of the variation explained—see Appendix Table A.1 in Tucker (2026).

across occupations.⁵ Fourth, there is heterogeneity in the exposure-adoption relationship; in this case, the heterogeneity is across workers and tasks.

Research with European data has also investigated the link between AI exposure and adoption. Using German survey data on workers linked to establishment-level data, Lindenthal et al. (2026) find that the exposure measure of Eloundou et al. (2024) explains only about 30% of the variation in actual worker-reported adoption at the occupational level, whereas a measure of AI adoption based on AI’s comparative advantage relative to a specific worker explains almost 60% of the cross-occupation variation in AI use. Using matched worker-firm data from Denmark, Pulito et al. (2026) similarly find that firm-level AI exposure measures based on the 2015 workforce composition do not strongly predict AI adoption at the firm level in 2023 and 2024.

Recent work also shows that LLM-based exposure measures can be sensitive to the model used to construct them, with resulting differences in occupational exposure scores and downstream estimates (Yin et al., 2026). Our analysis complements this work by comparing a broad set of exposure measures with direct firm-level AI adoption.

The paper proceeds as follows. Section 2 lays out the conceptual framework that motivates the empirical analysis. The data and linkages are described in Section 3. Section 4 summarizes the findings, followed by a discussion and concluding remarks in Section 5.

2 Conceptual framework

Let E_f denote the measured AI exposure of firm f and A_f be a binary indicator of its actual adoption or use of AI. Exposure captures the technological applicability of AI to the firm’s activities, whereas adoption, a decision variable, captures realized use of the technology. The two are therefore distinct economic objects, even though exposure may be an important

⁵Since occupations only account for about 20% of variation in adoption, exposure measures account for a relatively small fraction of the variation in worker-level adoption.

determinant of adoption. Moreover, exposure, adoption, and other firm characteristics (X_f) may each be related to firm outcomes, Y_f , through distinct economic channels. Figure 1 provides an illustration of these firm-level relationships.⁶

A simplified formulation of the adoption decision is

$$A_f = \mathbf{1} \{B(E_f, X_f) - C_f + \eta_f > 0\}, \quad (1)$$

where $B(\cdot)$ denotes expected benefits from adoption, C_f captures implementation and adjustment costs, and η_f represents unobserved firm-specific determinants of adoption.⁷ Equation (1) makes clear that predetermined exposure can be viewed as one of the state variables in the optimization problem of a firm where adoption is the decision variable. However, it does not imply a one-to-one, or even a monotonic, relationship between measured exposure and adoption without additional restrictions.

Predetermined exposure is not necessarily exogenous, as Figure 1 clarifies. Firm-level exposure measures are often constructed from pre-existing occupational or task composition, which may reflect persistent differences in existing firm characteristics. These characteristics may independently affect both subsequent AI adoption and economic outcomes.

Most widely used AI exposure measures are not originally constructed at the firm level. They typically map AI capabilities to tasks, abilities, or occupations and can be aggregated to industries, geographic areas, or firms using employment composition. For example, firm-level workforce exposure can be constructed as an employment-weighted average of occupational exposure scores. Such a measure provides firm-specific variation, but it remains fundamentally a measure of the applicability of AI to worker tasks rather than a direct

⁶The figure is a simplified representation of the underlying relationships for illustrative purposes. For instance, firm characteristics are depicted as exogenously given, though some of the characteristics, and exposure, could change after adoption. Moreover, factors like uncertainty and anticipation can affect adoption directly, in addition to outcomes. These effects, and any dynamic or intertemporal considerations, are omitted for simplicity.

⁷These could include idiosyncratic beliefs and uncertainty about the effects of AI, as well as other unobserved factors, such as managerial talent and organizational capabilities.

measure of the firm’s technological capabilities or AI use.

2.1 How closely does exposure relate to adoption?

Exposure and adoption may not coincide for at least two reasons. First, exposure captures technological susceptibility rather than the firm’s decision to adopt, whereas adoption also depends on other factors. A firm’s decision about whether to adopt AI is highly idiosyncratic, as many factors can play a part in the decision process, including the firm’s capacity to recognize this new technology (and its perceived suitability for firm activities); expected returns; implementation costs as balanced against cost savings (for example, from a lower wage bill); managerial and organizational capabilities; and access to complementary inputs, such as the skills and education of the workforce, as well as the software and digital infrastructure.

Second, measured exposure may capture only part of the firm’s relevant technological susceptibility. In particular, occupational or worker-task measures primarily characterize the labor side of production and may not fully capture AI exposure arising through, for example, software, equipment, and capital-embedded processes.⁸ These dimensions of exposure need not be highly correlated with worker-based exposure measures. Moreover, some exposure measures identify tasks that AI can affect without necessarily distinguishing whether AI substitutes for or complements the associated labor. This distinction may also matter for the adoption decision by altering expected benefits and costs.

A simple way to relate adoption to exposure is to estimate a linear probability model without assigning a causal interpretation

$$A_f = \alpha + \beta_{\text{raw}} E_f + \varepsilon_f. \tag{2}$$

⁸For example, a firm with software-controlled machinery likely has lower implementation costs of integrating AI into its processes than a firm whose machinery is operated manually.

This formulation examines whether more exposed firms are more likely to use AI. Similarly, a conditional specification is

$$A_f = \alpha + \beta_{\text{cond}} E_f + X_f' \gamma + \varepsilon_f, \quad (3)$$

where X_f contains firm characteristics, such as industry, size, and age. This examines how adoption is associated with exposure conditional on observed firm characteristics and within industries. Comparing β_{raw} with β_{cond} separates the component of the exposure–adoption relationship that operates across firm characteristics from the component that holds within cells defined by X_f .

2.2 Implications for exposure-based studies

AI exposure measures have increasingly been used to study labor-market and firm outcomes. For example, Acemoglu et al. (2022) construct establishment-level exposure from occupational task exposure and workforce composition and examine the associations between exposure and hiring, skill requirements, employment, and wages. Eisfeldt et al. (2026) aggregate occupational generative AI exposure to the firm level and examine the associations between exposure and firm valuation, profitability, and labor-market outcomes. Hampole et al. (2025) construct task-level AI exposure varying across firms, occupations, and time to study skill demand and within-firm labor demand and employment. At a different level of aggregation, Frank et al. (2025) relate occupational AI exposure to unemployment risk and job separations across occupations, states, and time. More recent work uses generative AI exposure to study post-ChatGPT labor-market outcomes: Brynjolfsson et al. (2026) examine employment and wages, documenting employment decline among younger workers in highly exposed occupations, while Tucker (2026) aggregates exposure to industry-state cells to study employment, hiring, separations, and earnings, with particular attention to early-career workers.

The underlying exposure measures are typically task- or occupation-based, while the unit

to which they are aggregated depends on the outcome and available data. Occupation-level exposure is naturally suited to studying how workers or jobs with different tasks experience technological change, but it does not directly speak to firms’ adoption decisions. Industry-level exposure summarizes the technological susceptibility of a sector but averages across firms with potentially wide underlying heterogeneity. Firm- or establishment-level exposure more closely matches the decision-making unit for adoption and permits within-industry comparisons, but it is generally constructed from the firm’s workforce composition and may not capture key dimensions of firms’ technology use.

These different aggregation levels can be represented by letting u denote the relevant empirical unit. Exposure-based studies often estimate specifications of the form

$$Y_{ut} = \alpha_u + \lambda_t + \theta(E_{ut_0} \times Post_t) + v_{ut}, \quad (4)$$

where Y_{ut} is an outcome and E_{ut_0} is predetermined exposure in some period t_0 . The interpretation of θ depends on the unit of analysis (u). At the occupation level, it captures differential outcomes for more exposed jobs; at the industry level, it captures differential outcomes for more exposed sectors; and at the firm or establishment level, it captures differential outcomes for organizations with more exposed workforces.

A weak conditional exposure–adoption relationship does not imply that exposure measures are uninformative. Rather, it clarifies what they measure or capture. If exposure only imperfectly predicts adoption, θ in equation (4) cannot generally be interpreted as the effect of firms actually adopting AI. It instead captures the differential response to the AI technology shock of units that were more technologically susceptible to AI prior to the shock. This is an economically meaningful concept, but it can differ from the effect of realized adoption, which may operate through additional channels.

The distinction also matters for interpreting the magnitude of exposure-based estimates. Suppose that the true underlying relationship between an outcome and adoption is such that

adoption is the only variable relevant for the outcome

$$Y_f = \alpha + \tau A_f + v_f, \quad (5)$$

and exposure is related to the outcome only through adoption as in (2). Let $\theta_{Y,E}$ denote the population coefficient from a regression of Y_f on E_f . One can then write

$$\theta_{Y,E} = \tau \frac{\text{Cov}(E_f, A_f)}{\text{Var}(E_f)} + \frac{\text{Cov}(E_f, v_f)}{\text{Var}(E_f)}. \quad (6)$$

Note that the population coefficient in (2) is $\beta_{\text{raw}} = \frac{\text{Cov}(E_f, A_f)}{\text{Var}(E_f)}$. If we further assume $\text{Cov}(E_f, v_f) = 0$, equation (6) reduces to

$$\theta_{Y,E} = \tau \beta_{\text{raw}}. \quad (7)$$

Thus, under the imposed simplifications, even a substantial effect of adoption can generate a small exposure–outcome coefficient when exposure is only weakly associated with adoption.

More generally, the simplifying assumptions underlying equation (7) need not hold. As shown in Figure 1, exposure may be correlated with other determinants of outcomes, and it may affect outcomes through channels other than measured firm adoption. A small exposure coefficient therefore does not imply that the adoption–outcome relationship is weak. Conversely, a large exposure coefficient does not imply a strong adoption–outcome connection, because it may reflect a strong exposure–adoption mapping (large β_{raw}), channels associated with exposure that operate outside measured adoption, or a correlation between exposure and other determinants of outcomes (large $\text{Cov}(E_f, v_f)$). For example, a highly exposed firm may alter the composition of its workforce in anticipation of its future adoption and reliance on AI, in response to rivals’ adoption of AI, or to postpone or reduce hiring while evaluating AI. Exposure could thus have predictive power for outcomes because it measures the firm’s incentives to adjust, even in the case of limited adoption.

Finally, the observed relationship between exposure and adoption may also reflect measurement error in adoption, a binary dependent variable. If some firms incorrectly report their adoption status in survey data (i.e., misclassification), the estimated relationship between exposure and adoption may differ from the relationship with true AI use.⁹ A weak observed exposure–adoption relationship may therefore reflect a gap between exposure and adoption, measurement limitations in exposure, measurement error in adoption, or some combination of these factors. The estimates are therefore best interpreted as characterizing the empirical association between exposure measures and reported measures of AI use, rather than as identifying a structural and causal relationship.

3 Data

We describe the data used in the analysis starting with occupational exposure measures followed by adoption measures and how we link the two sets of measures at the firm level.

3.1 Occupational composition

Establishment-level occupational employment data come from the BLS Occupational Employment and Wage Statistics (OEWS) program. This program surveys approximately 200,000 establishments every May and November and respondents report the number of workers and their wages in roughly 830 detailed occupations defined in the Standard Occupational Classification (SOC) system. Since AI adoption potentially changes an employer’s occupational needs, we use an establishment’s reported staffing pattern before the widespread

⁹While classical measurement error in a continuous dependent variable does not bias OLS coefficient estimates, it increases residual variance and reduces precision. When the dependent variable is binary, misclassification can bias the estimates—see Meyer and Mittag (2017) for an analysis.

adoption of generative AI technology to construct measures of its AI exposure. Therefore, we restrict our attention to responding units surveyed from May 2018 to November 2020.¹⁰

3.2 AI exposure measures

Researchers estimate how occupations, firms, or industries may be affected by AI and advanced computing technologies by analyzing their relevance for “tasks,” “work activities,” or “abilities” documented in the Occupational Information Network (O*NET), a comprehensive and regularly updated database that reports required skills, education, job tasks, and work conditions for detailed occupations.¹¹ Measures of potential exposure to large language models (LLMs) can be aggregated to O*NET occupations and then mapped to Standard Occupational Classification (SOC) occupational categories, which differ slightly from O*NET occupations. Using occupational employment shares from the OEWS matrices, these occupation-level measures have been aggregated to construct firm- and industry-level measures of LLM exposure. Several AI-exposure scoring systems exist. These can be broadly divided into two groups—those based on ex ante task content and those based on some form of observed AI usage. The latter might capture the relationship between exposure and adoption more directly. Below, we summarize the ones used in our analyses.

The literature frequently cites two occupation-level measures from two papers for assessing theoretical exposure to LLMs: those developed by Eloundou et al. (2024) and one by Felten et al. (2023). Eloundou et al. (2024) develop six measures to assess whether LLMs can assist workers in performing their existing tasks more quickly, either alone or when integrated with other software applications and systems.¹² Their measures of exposure at the O*NET task level rely on three metrics—referred to as alpha, beta, gamma—each rated by either human evaluators or OpenAI’s GPT-4. These metrics assess: (1) whether tasks can

¹⁰OpenAI released ChatGPT-3.5 to the general public in November 2022. However, enterprise AI tools were in use in early 2021 (Dohmke, 2022; Friedman, 2021).

¹¹U.S. Department of Labor, Employment and Training Administration 2026.

¹²The LLMs are evaluated based on their capabilities in text, code, and image generation in mid-2023.

be completed at least 50% faster using LLMs alone or with a simple interface, (2) whether tasks can be completed at least 50% faster using LLMs alone or in combination with software, such as domain-specific software, with a 50% weight on tasks that need associated software, and (3) whether tasks can be completed at least 50% faster using LLMs alone or in combination with software, where they place equal weight on tasks that do or do not need additional software to reduce the time necessary to complete a task by 50%.¹³

Eloundou et al. (2024) show that GPT-4-rated and human-evaluated metrics are highly correlated, and use the second metric described above that was rated by human evaluators to draw conclusions about the potential impact of generative AI on jobs. They find that over 18.5% of workers in the 2020-2022 OEWS have at least 50% of their tasks potentially exposed to LLMs and that higher-wage occupations are more exposed than lower-wage occupations. Their GPT-4 based measure characterizes LLM applications that could be easily implementable in businesses' workflows at the time of the evaluation, and thus it may be more highly correlated with realized adoption in the short run. They aggregate the task exposures to the O*NET occupation level and then to the 2018 SOC occupations. Their measures can be interpreted as the percentage of tasks in an occupation that can be done 50% faster with LLMs.

Felten et al. (2021) construct occupational, industry, and geographic measures of exposure to a combination of 10 types of AI applications. To cover generative AI, Felten et al. (2023) further construct exposure to LLMs for text generation and image generation separately. Their measures are built by linking AI functions to 52 distinct human abilities required to perform jobs rather than tasks. Their occupation-level exposure measures are based on the 2010 SOC, which matches our OEWS data less well than the 2018 SOC does. Felten et al.'s measures are standardized at the occupation level with a mean of zero and standard deviation of one.

¹³Eloundou et al. (2024) label the share of tasks that can be done by LLMs alone as E1 and the share of tasks that can be done by LLMs in conjunction with other software applications and systems as E2.

Eisfeldt et al. (2026) create separate measures for whether the ‘core tasks’ of an occupation versus its less-central ‘supplemental tasks’ are likely to be affected by LLM capabilities. They theorize that AI services that displace the core tasks of an occupation are more likely to reduce employment in that occupation (‘labor substitution’), whereas the automation of supplemental tasks complement the core capabilities of people in an occupation, and would improve the productivity of those workers without displacing them. They apply their measures to U.S. publicly traded firms and find that the labor substitution effect of automating core occupation tasks has a greater effect on employment. Their occupation-level exposure measures also are based on the 2010 SOC.

Unlike the earlier authors, Massenkoff and McCrory (2026) construct a measure of AI displacement (observed exposure) that is based partly on data on actual usage of an LLM (Anthropic’s Claude) for work-related tasks as developed by Handa et al. (2025). Using tasks exposed to AI according to Eloundou et al. (2024), they label the tasks as exposed if they see sufficient work-related Claude usage traffic. They use this data to infer, for instance, that while, in principle, most mathematical tasks could be assigned to AI systems, much less than half now are. Their measure reflects initial steps toward implementation and automation, and places more weight on automating uses. They also use information from O*NET to adjust for the importance of tasks within occupations. Their measure is best viewed as a usage-adjusted exposure measure that lies between an exposure measure and direct measures of firm or worker adoption. Their occupation-level exposure measures are based on the 2018 SOC.

Tomlinson et al. (2025) analyze user conversations with Microsoft Copilot to measure AI applicability to occupations. They map O*NET 29.0 Database intermediate work activities assisted and performed successively by Copilot in each conversation to multiple occupations through tasks. Their occupation-level AI applicability score is based on the 2018 SOC and incorporates task importance and relevance as weights. It allows relative comparisons across occupations as opposed to measuring the share of work activities that can be done with AI.

Using their measure, they predict which occupations will delegate tasks to AI and which will use AI to assist existing workflows. They find that AI-assisted work is most common for information work.

Because establishments and exposure measures are coded under different SOC systems, we anchor exposure measures to a common occupational classification. We use the occupation level in the May 2019 OEWS, for which occupations from both the 2010 and 2018 SOC systems can be mapped consistently.¹⁴ For exposure measures defined at a more detailed level, we aggregate them to this common occupation level using OEWS employment shares. Measures based on the 2010 SOC use May 2018 employment shares, while measures based on the 2018 SOC use May 2021 shares, because these years provide full employment coverage under the respective SOC systems. For example, because the Eloundou et al. (2024) measures are based on the 2018 SOC, we aggregate them using May 2021 OEWS employment shares. The resulting exposure value is then assigned to both the 2010 and 2018 SOC occupations that map to the same May 2019 published occupation.¹⁵

Table 1 presents the definitions, sources, and labels of the exposure measures. These occupation-level measures capture different dimensions of AI exposure and exhibit substantial variation in their pairwise correlations, which range from 0.26 to 0.98 across occupations in the May 2019 OEWS (see Table A1). Felten et al.’s original AI exposure measure (aioe) is nearly perfectly correlated with their language modeling measure (lm_aioe; $\rho = 0.98$), and somewhat less strongly correlated with their image generation measure (ig_aioe; $\rho = 0.85$). Eisefeldt et al.’s core-task measure (gec) is more strongly correlated with the other measures than their supplemental-task measure (ges). Eloundou et al.’s beta and gamma measures, whether rated by humans or GPT (drb, hrb, drg, hrg), correlate more strongly with lan-

¹⁴Only the May 2018 OEWS observations in our sample use the 2010 SOC system. Subsequent observations use the 2018 SOC.

¹⁵Relatively few exposure measures require aggregation from more detailed occupations. For the 2018 SOC, for example, 786 detailed occupations in the OEWS have exposure data, but only 59 are aggregated into 24 published occupations. In addition, exposure values are generally similar across the detailed occupations within these aggregates.

guage modeling AIOE ($\rho > 0.80$) than do their alpha measures (dra, hra). Massenkoff and McCrory’s observed exposure measure (oe) correlates most strongly with Eisfeldt et al.’s total exposure measure ($\rho = 0.65$), while Tomlinson’s AI applicability score (aiapp) correlates most strongly with Eloundou et al.’s human-rated beta measure (hrb; $\rho = 0.75$). The two measures highlighted in our empirical analysis, ‘drb’ and ‘oe’, have a correlation of 0.63.

O*NET occupations do not perfectly cover SOC occupations partly because the SOC includes more “not elsewhere classified” occupations.¹⁶ There are a variety of ways to handle that small gap, which we face for any of these classifications. In this paper, we treat these residual occupations as missing. We derive estimates of each establishment’s AI exposure by weighting each non-missing occupational exposure measure by the occupation’s share of establishment employment. For most of the measures (Eloundou et al., 2024; Eisfeldt et al., 2026; Massenkoff and McCrory, 2026), the interpretation of the measure is the share of the establishment’s tasks that are potentially exposed. Only Felten et al.’s measures are standardized normal exposure measures. For our establishment-level analysis, we use a standard normal cumulative distribution function to transform the standardized measures to a near-uniform distribution on the interval $[0,1]$.

3.3 AI adoption measures

Our measures of firm-level AI adoption come from the Census Bureau’s Business Trends and Outlook Survey (BTOS), a nationally representative biweekly survey of nonfarm employer businesses. The BTOS annual sample contains approximately 1.2 million businesses, divided into six collection panels of roughly 200,000 businesses each. A collection cycle lasts 12 weeks, with one panel surveyed every two weeks. Depending on the sample year, the BTOS conducts four or five cycles, so each sampled business reports once every 12 weeks.¹⁷

¹⁶For example, O*NET covers “All other managers,” SOC 11-999, but not “all other sales and related workers,” SOC 41-9099.

¹⁷The timing of collection, reference periods, and publication is important for interpreting the data. Responses are collected during the two weeks following the reference period and are published approximately

The BTOS includes a set of core questions that are asked throughout the sample year. Beginning in September 2023, these questions included two measures of AI use: one concerning current use and another concerning expected use. The current-use question asks whether the business used AI during the previous two weeks, while the future-use question asks whether it expects to use AI during the next six months. The two questions are otherwise phrased similarly. The original current-use question asked: “Between MMM DD–MMM DD, did this business use Artificial Intelligence (AI) in producing goods or services? (Examples of AI: machine learning, natural language processing, virtual agents, voice recognition, etc.)” Respondents could select “Yes,” “No,” or “Don’t Know.” The future-use question used the same structure but asked whether the business expected to be using AI during the next six months. The phrase “in producing goods or services” was intended to capture substantive integration of AI into business operations rather than incidental use.¹⁸

As AI technologies diffused rapidly, particularly following advances in generative AI, researchers and other stakeholders raised concerns that the phrase “in producing goods or services” might be too restrictive and could understate business use of AI. Generative AI, in particular, enabled applications across a broad range of business functions that respondents might not naturally view as part of “producing goods or services.” Cognitive testing similarly indicated that broader wording would help respondents recognize and report AI use across a wider range of functions and tasks. In response, the Census Bureau revised the BTOS core AI questions in November 2025, replacing “in producing goods or services” with “in any of its business functions.” The revised wording therefore captures a broader range of AI use, including some applications that may not have been reported under the original phrasing.¹⁹

two weeks later. Much of the analysis in this paper uses data from the “revised AI core question,” first collected during November 17–30, 2025, referring to November 3–16, 2025, and first published on December 4, 2025. See Bonney et al. (2026), footnote 5.

¹⁸As Bonney et al. (2026) note, this wording closely followed a question used in the 2019 Annual Business Survey to measure business use of AI.

¹⁹The effect of the wording change cannot be estimated precisely because it coincided with the launch of the second AI supplement and followed a six-week federal government shutdown during which the BTOS was

The BTOS also periodically fields topical supplements that run for a single 12-week collection cycle. Two of the three supplements fielded to date have focused on AI. The first AI supplement was conducted from December 2023 through February 2024 (see Bonney et al. (2024)). Our main analysis uses the second AI supplement, conducted from November 2025 through February 2026, together with the revised core AI question. For comparison, we also report results using the original pre-revision core question and document how the measured exposure–adoption relationship differs across the two formulations.

3.4 Data linkage and analysis sample

We create our analysis sample by linking both OEWS respondents and BTOS respondents to the Census Bureau’s Longitudinal Business Database (LBD), a longitudinally consistent dataset tracking the universe of private nonfarm businesses in the U.S. The LBD is based on the Census Bureau’s Business Register (BR), which is also the sampling frame used for the BTOS. As a result, we can match the BTOS directly to the LBD using common BR identifiers. Since the BTOS is a firm-level survey, we correspondingly match respondents to the LBD at the firm level.

To link the OEWS to the LBD, we use two different matching methods and keep linked records found using either one. First, we match the OEWS to the Longitudinal Employer Household Dynamics (LEHD) Infrastructure, a data system created from the same state-level data used by the BLS to sample the OEWS, using state business identifiers (SEIN). From the LEHD database, we obtain a federal business identifier (EIN) which then matches to the LBD. This method has the advantage of using the curated state-to-federal business identifier crosswalk created by the LEHD production system. The OEWS also has EINs and we use these to do a second match directly to the LBD, which allows us to expand our matches. We match each year of sampled OEWS establishments to the corresponding LBD

not collected. Bonney et al. (2026) provide an approximate assessment of the resulting change in measured adoption.

year for the years 2018–2020. Since this match uses EINs, it does not match a specific OEWS establishment to a specific LBD establishment. Rather, it matches an OEWS establishment to an LBD firm.

Once the matching is complete on both sides, we link OEWS-LBD firms to BTOS-LBD firms using the LBD firm identifier. This creates a sample of firms that operated in the 2018–2020 period and in the 2023–2025 period. Because we use the LBD firm identifier to link, the firm must have used the same EIN in both periods, ruling out any firm reorganizations that changed the EIN. This conservative match has the advantage of only including firms in our sample that did not substantially change in organizational structure between when they were sampled by the OEWS and when they were sampled by the BTOS. In the future, we will expand these matching efforts to use other identifiers to match across firm reorganizations. However, we do not expect this will substantially alter our sample because in recent Census business data, less than 1% of single-unit firms match using alternative linking methods (Chow et al., 2021).

After completing the firm-level linking, we must reconcile the fact that the BTOS collects data on firm-level AI adoption and OEWS data pertains to establishments. Measuring AI exposure for a firm on the basis of the occupation distribution of one of its establishments could lead to substantial measurement error. In this paper, we handle this challenge by restricting to firms that operate in a single location, i.e., single-unit firms. For these firms, AI adoption and exposure are measured at the same level and our firm-level linking is sufficient to match the BTOS and OEWS at the analysis level. From the LBD we obtain firm characteristics such as industry, size, age, and location for the OEWS reference year.

The OEWS is designed around a three-year sample design that ensures each establishment unit is sampled no more than once in a three-year interval. As a result, repeated sampling of single-unit firms within our OEWS sample period (2018–2020) is unlikely. If we observe multiple OEWS responses for a single-unit firm, we use only information from the most recent response. Meanwhile, the BTOS creates a new sampling frame for each year,

but they also ensure that any sampled businesses are excluded from subsequent samples for at least one year. As a result, repeated sampling of single-unit firms across multiple BTOS sample years is also unlikely. If we observe the same single-unit firm across multiple BTOS sample years, we choose the BTOS response that is closest in time to the OEWS response.

In summary, our analysis sample is single-unit firms operating in at least one year between 2018 and 2020 and at least one year between 2023 and 2025 and linkable by EIN. To satisfy the data requirements of the regressions (2) and (3), a single-establishment firm is included in the sample if it responded to the core AI use question either before or after the Census Bureau changed its wording in 2025, and if all relevant exposure measures and firm characteristics are also observed. Using responses to the pre-revision question about AI use in producing goods or services yields a sample of approximately 57,000 firm-year observations. Using the revised adoption question about AI use in any business function results in a sample of about 15,500 firm-year observations.²⁰

The OEWS and BTOS are both designed to be representative of in-scope employer businesses, and with appropriate weighting, allow for the estimation of the underlying population parameters. However, because the sampling frames of these surveys are different, and the probability of survey response or record linkage may be related to firm characteristics, it is important to address the effects of sample selection. We use propensity score weights to this end: we estimate the relationship between the probability of being included in our sample and observed firm characteristics using yearly logistic regressions. The observables are industry-specific size bins based on corresponding employment and payroll quintiles calculated from the LBD. The logistic regressions contain industry fixed effects, industry-specific employment-size fixed effects, and common payroll-size fixed effects. The propensity score weights obtained from these regressions are designed to match the weighted within-industry size distribution in the universe of single-unit firms in the LBD.

²⁰ Although the post-revision sample is smaller, it is comparable to sample sizes in earlier studies that used OEWS-Census data; see for example, Blackwood et al. (2025).

Descriptive statistics for the variables used in the analysis are in Tables 2a and 2b. Average firm employment in our two samples is approximately 13, close to the average size values from the last two available years of the Business Dynamics Statistics for single-unit firms (9.7 and 9.9, respectively, for 2021 and 2022).²¹ We interpret these patterns as evidence that propensity score weighted statistics alleviate potential selection bias from our linkage procedure.

4 Results

We use our linked OEWS-BTOS data to directly test the empirical relationship between AI exposure and AI adoption by estimating the following weighted least squares regression at the firm level:

$$A_{ft} = \alpha + \beta E_{ft_0} + X'_{ft_0} \gamma + \varepsilon_{ft}. \quad (8)$$

A_{ft} is an indicator for whether firm f reported using AI in the BTOS in reference year t , and E_{ft_0} is a given measure of the firm’s AI exposure based on its occupational mix during the OEWS reference year t_0 . In some specifications, we also control for year effects and firm-level characteristics, such as industry, size, age, or state, through X_{ft_0} . This specification empirically implements the conceptual models from Equations (2) and (3). We define all controls using LBD data from the OEWS reference year t_0 . As a baseline, we use firm responses to the post-revision BTOS AI adoption questions in the 2025 reference year, and we weight our analysis sample to match the employment-weighted distribution of single-unit firms by industry and size. We have at most one observation per firm in our regression, so we do not cluster standard errors at a higher level. We instead report standard errors and 95% confidence intervals using heteroskedasticity robust standard errors.

²¹<https://www.census.gov/programs-surveys/ces/data/public-use-data/experimental-bds/bds-single-unit/data.html>

4.1 Economy-wide comparisons

We find that firms’ AI adoption decisions are robustly positively correlated with their degree of AI exposure; however, the strength of the relationship varies across exposure measures. Table 3 reports estimates from estimating Equation 8 using each of our two main exposure measures and an array of firm-level controls. In Column 1, we report the bivariate relationship between adoption and exposure without any additional controls. Our point estimates imply that firms with a 10% higher occupational exposure to AI, as rated by the GPT-4 beta measure from Eloundou et al. (2024), are 6.07 (s.e. 0.585) percentage points more likely to report using AI in any of their business functions. Our estimates using the exposure measure from Massenkoff and McCrory (2026) are similar, and imply that firms with a 10% higher observed exposure to AI are 6.48 (s.e. 0.686) percentage points more likely to report that they used AI in any of their business functions.

The relationship between adoption and exposure remains even after we account for granular differences in their industry, location, age, and size. In Column 2 of Table 3, we add fixed effects to control for the reference year (t_0) when the OEWS data was collected, which ranges between 2018 and 2020. Although, in principle, a firm’s occupational mix can vary due to differences from when these data were collected, in practice, we find that these differences are uncorrelated with firms’ adoption decisions. Our estimates in Column 2 are nearly identical to those in Column 1. In Columns 3, 4, and 5, we add progressively more detailed fixed effects for the firm’s industry at the 2-digit, 3-digit, and 4-digit NAICS level, respectively, while retaining the reference year fixed effects from Column 2. By including industry fixed effects, we control for average differences in exposure and adoption at the industry-level that are typically the basis for aggregate comparisons, and instead use our firm-level data to isolate the within-industry relationship between exposure and adoption.

We find that controlling for average industry-level differences reduces the correlation between adoption and exposure by 45–60%, but within-industry differences in firms’ occupational mix are nevertheless predictive of their AI adoption decisions. Adding additional

fixed effects for the firm’s age, size-class, and state in Column 6 actually *increases* the coefficient on exposure, which implies that omitting these variables exerts a negative bias on net. Our point estimates from Column 6 imply that even when we compare firms in the same NAICS-4 industry, size-class, age, and state, we still find that firms with a 10% higher occupational exposure to AI are 3.85 (s.e. 0.652) or 3.34 (s.e. 0.728) percentage points more likely to report using AI.

We note, however, that despite the robust positive correlation between adoption and exposure, the majority of variation in firms’ AI adoption decisions remains unexplained by their occupational AI exposure or their baseline firm characteristics. Table 3 also reports the adjusted R^2 of each regression, and in the specifications that include controls, the adjusted R^2 of the exposure measure after removing the contribution of the controls (within adjusted R^2). As reported in Column 1, firm-level occupational exposure to AI, as rated by the GPT-4 beta measure from Eloundou et al. (2024), explains approximately 5% of the overall variance in AI adoption, while the exposure measure from Massenkoff and McCrory (2026) explains approximately 3% of the overall variance in AI adoption. Adding simple fixed effects for the sector of the firm raises the share of overall variance explained to 7.6% for the regression using the GPT-4 beta exposure measure and 7.1% for the regression using the observed exposure measure. The majority of the share of variance explained is due to sector-level differences in average adoption (and exposure)—within-sector differences in firm exposure only explain 0.6–1.2% of the within-sector variance in AI adoption.

Adding more detailed industry fixed effects increases the share of the overall variance explained to approximately 12%, but the share of within-industry differences in AI adoption explained by within-industry differences in exposure is 0.3–0.7%. Further controlling for firm characteristics raises the adjusted R^2 by approximately 2 percentage points, and the share of residual adoption explained by residual exposure remains stable at 0.5–0.9%. These results highlight that industry-level factors and firm-specific characteristics can be useful predictors of firms’ AI adoption, but also that approximately 85% of the firm-level variation in AI

adoption remains unexplained.²²

We focus on two commonly used exposure measures in our analysis, but it is straightforward to extend our approach to consider other occupation-based exposure measures on a consistent basis. Figure 2 reports the coefficients from estimating Equation 8 using a range of additional AI exposure measures from Table 1. In each comparison, we use the same sample, adoption measure, regression weights, and firm-level occupational shares as do the regressions in Table 3, so differences arise solely from differences in occupational exposure to AI implied by each measure. Panel A of Figure 2 reports the coefficient estimates from a simple bivariate regression (the same specification as Column 1 of Table 3), while Panel B reports the coefficient estimates when controlling for year and NAICS-3 fixed effects (the same specification as Column 5 of Table 3).

Although the estimated relationship between adoption and exposure varies by measure, which is to be expected given that each exposure measure makes a different set of assumptions about each occupation’s exposure to AI, the results are qualitatively consistent. All 14 exposure measures are positively correlated with AI adoption, and these coefficients are all statistically significantly different from zero in bivariate regressions. When we include NAICS-3 fixed effects to isolate the within-industry relationship, the coefficients for all 14 measures remain positive, although we do not reject the null that the coefficient for one of the measures (the supplemental generative AI exposure measure from Eisfeldt et al. (2026)) is zero. Moreover, since each exposure measure ranges between 0 and 1, their relative magnitudes are informative. Panel A implies that a firm with a 10% higher exposure to AI is, on average, 4–12 percentage points more likely to report using AI. The coefficient tends

²²This low explanatory power is consistent with the finding of McElheran et al. (2024) from the 2018 Annual Business Survey (ABS) that a large set of firm and owner characteristics explains only a modest fraction of the variation in AI adoption. Ando et al. (2025) show that this finding applies more generally to the adoption of advanced technologies in the 2018 ABS. Acemoglu et al. (2025) reach a similar conclusion using the advanced technologies surveyed in the 2019 ABS. Relatedly, Bick et al. (2026) find exposure measures explain little of the variation in generative AI adoption at the worker-level and a modest amount of the variation in adoption by occupation.

to be higher for usage-based exposure measures, such as the AI applicability measure from Tomlinson et al. (2025) and the observed exposure measure from Massenkoff and McCrory (2026), and lower for the pre-generative AI exposure measure from Felten et al. (2021).

We note, however, that because the standard deviations of the measures differ considerably, the share of variance in AI adoption explained by each measure may not be monotonic with the magnitude of their coefficients.²³ Figure 3 compares standardized regression coefficients after scaling each exposure measure by its standard deviation within our employment-weighted firm sample. While standardizing coefficients would not change the sign or statistical significance of those estimates, it does change the interpretation of their magnitudes. Panel A implies that a firm with a one-standard-deviation higher exposure to AI is, on average, 4–11 percentage points more likely to report using AI. Moreover, the relative ranking of exposure measures in terms of coefficient magnitudes also changes, since the standard deviation of each measure generally differs. In contrast to Figure 2, we find that a one-standard-deviation increase in a firm’s usage-based exposure is correlated with among the smallest increases in AI adoption, whereas a standardized increase in pre-generative AI exposure predicts some of the largest increases in adoption. Controlling for year and industry fixed effects, a one-standard-deviation increase in exposure is associated with a 1–8 percentage points higher adoption probability.

We also assess the degree to which our main conclusions depend on methodological choices around sample selection and weighting. Table A2 repeats the analysis in Table 3 on the subsample of firms with at least 5 employees in the OEWS reference year. We also impose that the employment recorded in the LBD and the OEWS in the reference year differ by no more than 50% to ensure that the entities in the two frames plausibly correspond to the same establishment. Table A3 also repeats the analysis in Table 3, but weights the sample to match the firm-weighted distribution of single-unit firms rather than

²³See Table 2a and 2b for the standard deviation and the coefficient of variation for each exposure measure.

the employment-weighted distribution. Our conclusions remain qualitatively similar in both cases. The coefficient estimates are generally slightly larger in the subsample that excludes the smallest and potentially mismatched firms, and the coefficient estimates are smaller in the firm-weighted case. Both sets of results imply that the correlation between adoption and exposure may systematically differ by firm size; we systematically evaluate these differences in the next subsection.

Finally, although our preferred measure of AI adoption is the post-revision measure from BTOS, we also explore the relationship between exposure and the pre-revision measure from BTOS. The pre-revision measure asks whether firms used AI in producing goods or services rather than for any business functions. As a result, it may capture a distinct form of AI usage. Since the language was in use from 2023–2025, it is also informative about the earlier periods of AI adoption. As Table 2b notes, the (employment-weighted) average AI adoption rate is 9% in the pre-revision sample as compared with 27% in the post-revision sample. Despite these potential differences, Table A4 shows that the relationship between the pre-revision AI adoption and our two main exposure measures is qualitatively similar to our results using the post-revision AI adoption measure. Moreover, Figure A1 extends this comparison to all 14 exposure measures, and we again find that the correlation between AI adoption and exposure remains positive with varying degrees of strength (with the exception of the generative-AI supplemental exposure measure after including NAICS-3 fixed effects).

4.2 Differences by firm size and sector

An advantage of our matched firm-level data is that we can use a consistent framework to unpack the degree to which adoption behavior differs from exposure for different groups of firms. We focus on two firm characteristics that are important determinants of AI adoption in our sample: firm size and sector.

To understand differences by firm size, we divide firms into five size classes based on their total employment in the OEWS reference year, and we separately estimate the relationship

between AI adoption and exposure for each size class. Table 4 reports estimates for the bivariate regression coefficient, as well as average adoption and exposure measures for each of our two main exposure measures. Consistent with our firm-weighted results, we find that the relationship between AI adoption and exposure is 40–60% larger for firms with at least 20 employees as opposed to firms with fewer than 5 employees — although our estimates for the largest size class (100+ employees) are less precise, so those differences are not statistically significantly different from the other estimates. These larger coefficients help reconcile the fact that large firms tend to have *lower* levels of AI exposure but higher levels of AI adoption.

We also unpack sector-level differences by separately reporting the average adoption, exposure, and the bivariate coefficient between the two measures by sector in Table 5. Since the number of firms in some sectors may be small, we restrict the sample in each sector to our “high-confidence matches” (firms with at least 5 employees and where the employment recorded in the LBD and the OEWS in the reference year differ by no more than 50%) to ensure that our sectoral differences are not determined by the relative proportion of very small or potentially mismatched firms.²⁴

Table 5 highlights the distinction between adoption and exposure at the sector level.²⁵ Sectors with a high degree of AI adoption do not necessarily have a high correlation between adoption and exposure. Information (NAICS 51) has the highest level of average AI adoption (58.2%), and there is also a strong relationship between AI adoption and AI exposure in this sector; differences in AI exposure due to the occupational mix of firms in this sector can explain up to 16% of the within-sector variation in AI adoption (7% when using the observed exposure measure). On the other hand, in Education Services (NAICS 61), AI adoption is also high (37.1%), yet there is no statistically significant relationship between AI adoption

²⁴We also exclude the following sectors from this analysis because they have relatively few firms that responded to both the BTOS and the OEWS, and the estimates from those sectors may be unstable: Agriculture, Forestry, Fishing and Hunting (NAICS 11), Mining, Quarrying, and Oil and Gas Extraction (NAICS 21), Utilities (NAICS 22), and Management of Companies and Enterprises (NAICS 55).

²⁵These patterns may be easier to see in a figure, so we also present these results in Figure A2.

and the AI exposure of firms in this sector. The same contrast exists among sectors with low AI adoption. While AI adoption is strongly related to AI exposure in Manufacturing (NAICS 31–33) and Wholesale Trade (NAICS 42), the two measures are essentially uncorrelated in Accommodation and Food Services (NAICS 72).

Finally, our framework provides a simple benchmark for the sectors that are either more or less likely to adopt AI relative to their level of exposure. Specifically, we use the economy-wide bivariate regression coefficients from our baseline bivariate regression (Table 3, Column 1) to map each sector’s average AI exposure (as reported in Table 5) to a predicted AI adoption rate. We plot these actual vs exposure–predicted sectoral adoption rates in Figure 4. These predicted adoption rates are descriptive (in the spirit of a Kitagawa-Oaxaca-Blinder decomposition) rather than estimates of causal effects.²⁶ But since our adoption and exposure measures are drawn from the same sample of firms, this exercise provides a useful opportunity to quantitatively compare the sectoral differences in adoption versus exposure without the concern that the two measures are drawn from different sets of firms.

The Information (NAICS 51) and the Professional, Scientific, and Technical Services (NAICS 54) sectors have the highest (employment-weighted) rates of AI adoption and some of the highest rates of AI exposure.²⁷ The Finance and Insurance sector (NAICS 52) has a similar level of AI exposure as the Information and Professional, Scientific, and Technical Services sectors, but a substantially lower AI adoption rate. Our quantification shows that the AI adoption rate in Finance and Insurance is actually typical given its AI exposure, whereas the rate of AI adoption in the Information and Professional, Scientific, and Technical Services sectors is 8–22 percentage points higher than what would have been predicted based

²⁶Our exercise is similar to the exercise in Bick et al. (2026), which uses worker-level data to identify occupations where exposure measures either over- or under-predict AI adoption.

²⁷These two sectors have consistently been found to be the highest adopters of AI based on the analysis of the BTOS AI supplements, followed by Education Services (NAICS 61) and Finance and Insurance (NAICS 52) — see Bonney et al. (2024) and Bonney et al. (2026).

on their exposure alone.²⁸

5 Discussion and concluding remarks

Our findings indicate that exposure is positively related to adoption but explains a limited fraction of the variation in adoption. This result is important for interpreting studies that relate AI exposure to firm outcomes. Exposure measures identify where AI is technologically relevant to existing tasks or occupations, an economically meaningful object in its own right, but our results show that the exposure-adoption mapping is not close to one-to-one. Exposure–outcome relationships may operate through adoption, as realized AI use affects productivity, labor input, output quality, and complementary investment, but they may also reflect other channels. Firms may, for example, adjust hiring, skill requirements, investment, or organization in anticipation of future adoption or in response to uncertainty about AI’s effects on their own activities and those of their competitors.

This distinction also matters for interpreting the magnitude of the exposure–outcome relationships. Exposure imperfectly predicts adoption, and some exposure measures are more weakly correlated with adoption than others. A small coefficient on exposure need not imply a weak relationship between realized AI use and the outcome; conversely, a large coefficient may reflect adoption, anticipatory adjustments, or characteristics correlated with exposure. For example, if hiring declines more in highly exposed firms or industries, the relationship may reflect reduced hiring by firms already using AI, anticipatory adjustments by exposed nonadopters, or both. A highly exposed firm may postpone filling a vacancy while testing AI, redesigning the job, or learning how the technology will affect future labor needs. Such behavior can generate an exposure–hiring relationship even before realized AI use. Thus, without additional assumptions or evidence on adoption, an exposure–outcome relation-

²⁸Using worker data from a household survey, Bick et al. (2026) find that measures of exposure overpredict generative AI usage in clerical and administrative occupations and underpredict usage in occupations with a relatively high degree of autonomy including computer and research information scientists.

ship cannot generally be interpreted as the relationship between AI use and the outcome. Our analysis does not seek to identify how much of a given exposure–outcome relationship operates through adoption versus other channels. Instead, it complements exposure-based studies by documenting how technological susceptibility maps into realized use and sharpens the interpretation of exposure-based estimates.

There may be several mechanisms behind the exposure–adoption gap. Occupational exposure captures opportunities for AI to substitute for or augment tasks, while realizing those opportunities may depend on firms’ capacity to recognize, assimilate, and apply new technologies—what Cohen and Levinthal (1990) term absorptive capacity.²⁹ Relevant factors may include prior technological experience, digital infrastructure, complementary software and data, worker skills and education, and managerial/organizational capabilities.³⁰ Consistent with these factors, Bonney et al. (2026) document systematic relationships between the breadth of AI integration, operational investments, and firm outcomes. We could leverage additional information about firms’ perceptions about barriers to AI use from BTOS supplement questions (especially the question concerning reasons a firm does not plan to use AI in the near future). We also plan to use additional micro-level datasets to help us understand the role of factors that are relevant for adoption, such as management practices, intangible capital, prior investments in capital, and worker characteristics. For example, we plan to examine worker characteristics by linking in data from the LEHD dataset.

The firm-size and sector results are also consistent with a role for complementary factors. The exposure–adoption relationship is generally stronger among larger single-unit firms (though not uniformly different statistically across all firm size classes), even though aver-

²⁹These considerations also connect to the broader diffusion literature, which emphasizes systematic differences among adopters of new technologies (Rogers, 1962).

³⁰Some occupational AI exposure measures can incorporate the role of education indirectly. Vasilaukas and Horrigan (2026) examine occupational AI exposure together with BLS employment projections and show that among highly exposed occupations, projected employment growth is concentrated in occupations requiring a bachelor’s degree or more, while occupations with lower educational requirements have weaker or negative projected growth.

age occupational exposure does not increase monotonically with size. Larger firms may be better able to surmount implementation barriers or possess more of the complementary inputs needed to convert exposure into use. Similarly, occupational composition alone does not account for the large differences in adoption across sectors. In some high-adoption sectors, such as Information and Professional Services, the relationship is strong, but in other high-adoption sectors, such as Education Services, the relationship is weak. The observed exposure–adoption gap also helps with the interpretation of the current relatively low rate of AI adoption across firms:³¹ technological applicability often does not translate into realized use. The exposure–adoption link may become tighter as AI becomes cheaper, more capable, and easier to integrate. Highly exposed nonadopters may adopt as implementation costs fall, while the application of AI in new tasks may alter which exposure measures best predict use.

The firm-level exposure–adoption relationship is complementary to the corresponding worker-level relationship. Bonney et al. (2026) show that AI diffusion can involve both firm-directed and worker-driven use, highlighting the relevance of top-down and bottom-up adoption channels. Firm adoption takes into account business-wide expected benefits, implementation costs, complementary investments, and managerial capabilities, whereas worker adoption is more directly tied to the gains and costs of using AI for particular tasks. However, the two decisions are interdependent. Firm policies can facilitate worker use, and worker experimentation can encourage broader firm adoption. Exposure may therefore translate into adoption differently at the firm and worker levels, even though the two are closely connected. Future work can further explore firm-level and worker-level AI usage using matched data on firm and worker adoption. Data collected on worker AI use assessed at the firm level in the second BTOS AI supplement and individuals’ work-related AI use measured in the Census Bureau’s Household Trends and Outlook Pulse Survey (HTOPS) can be useful

³¹The BTOS-based estimate of the adoption rate among U.S. firms was about 22% as of this writing.

in this direction.

We reiterate some limitations to qualify the interpretation of the results. The analysis is descriptive and does not identify a causal effect of exposure on adoption. Occupational exposure measures may miss technological exposure to AI arising, for instance, from software, equipment, and capital-embedded processes. Reported AI use itself may be subject to misclassification when respondents differ in what they recognize as AI, and the binary measure of adoption does not capture the intensity of use. Nevertheless, the similarity of the qualitative results for pre- and post-revision adoption measures from BTOS is reassuring. Finally, the restriction of the analysis sample to single-unit firms limits the applicability of the results, because multi-unit firms tend to be much larger and heavier AI adopters at the same time (Bonney et al., 2026). In future work, we plan to expand our sample to include multi-unit firms.

This analysis is the first in a planned series of studies from the BLS–Census Collaborative Micro-productivity Project aimed ultimately at understanding how AI adoption relates to firm performance, especially productivity. In the framework of Figure 1, this paper focuses on the empirical mapping from exposure to adoption. The next step is to connect exposure and adoption to economic outcomes by leveraging the complementary microdata available at the BLS and Census Bureau. More broadly, distinguishing technological exposure from realized use provides a foundation for understanding how the diffusion of AI translates into changes in firm performance and productivity.

References

- Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022. doi: 10.1086/718327.
- Daron Acemoglu, Gary W. Anderson, David N. Beede, Catherine Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas. Automation and the workforce: A firm-level view from the 2019 annual business survey. In Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner, editors, *Technology, Productivity, and Economic Growth*, pages 13–55. University of Chicago Press, Chicago, 2025. URL <https://www.nber.org/books-and-chapters/technology-productivity-and-economic-growth/automation-and-workforce-firm-level-view-2019-annual-business-survey>.
- Yoshiki Ando, Emin Dinlersoz, Jeremy Greenwood, and Ruben Piazzesi. Technifying ventures. NBER Working Paper 33993, National Bureau of Economic Research, July 2025. URL <https://www.nber.org/papers/w33993>.
- Alexander Bick, Adam Blandin, David Deming, and Tyler Schumacher. What does generative AI do? NBER Working Paper 35677, NBER, 2026.
- Glenn J. Blackwood, Cindy Cunningham, Matthew Dey, Lucia Foster, Cheryl Grim, John Haltiwanger, Rachel Nesbit, Sabrina Pabilonia, Jay Stewart, Cody Tuttle, and Zoltan Wolf. Job tasks, worker skills, and productivity. Working Paper CES-25-63, Center for Economic Studies, US Census Bureau, September 2025.
- 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, April 2024.
- Kathryn Bonney, Cory L. Breaux, Emin Dinlersoz, Lucia S. Foster, John C. Haltiwanger, and Aditya A. Pande. The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks. NBER Working Paper 35141, National Bureau of Economic Research, April 2026.
- Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab, August 2026. Revised version; originally published in 2025.
- Melissa Chow, Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White. Redesigning the longitudinal business database. Working Papers 21-08, Center for Economic Studies, U.S. Census Bureau, May 2021. URL <https://ideas.repec.org/p/cen/wpaper/21-08.html>.
- Wesley M. Cohen and David A. Levinthal. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*, 35:128–152, 1990. doi: 10.2307/2393553.
- Thomas Dohmke. Github copilot is generally available to all developers, 2022. URL <https://github.blog/news-insights/product-news/github-copilot-is->

generally-available-to-all-developers/.

Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang. Generative AI and firm values. Working Paper 31222, National Bureau of Economic Research, January 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. doi: 10.1126/science.adj0998.

Edward W. Felten, Manav Raj, and Robert Seamans. A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57, 2018. doi: 10.1257/pandp.20181021.

Edward W. Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021. doi: 10.1002/smj.3286.

Edward W. Felten, Manav Raj, and Robert Seamans. How will language modelers like ChatGPT affect occupations and industries?, March 2023. URL <https://arxiv.org/abs/2303.01157>.

Morgan R. Frank, Yong-Yeol Ahn, and Esteban Moro. AI exposure predicts unemployment risk: A new approach to technology-driven job loss. *PNAS Nexus*, 4(4):pgaf107, 2025. doi: 10.1093/pnasnexus/pgaf107.

Nat Friedman. Introducing github copilot: Your AI pair programmer, 2021. URL <https://github.blog/news-insights/product-news/introducing-github-copilot-ai-pair-programmer/>.

Menaka Hampole, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research, February 2025. Revised September 2025.

Kunal Handa, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli. Which economic tasks are performed with ai? evidence from millions of claude conversations. Technical report, Anthropic, 2025.

Benjamin Labaschin, Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Extending “GPTs are GPTs” to firms. *AEA Papers and Proceedings*, 115:51–55, 2025. doi: 10.1257/pandp.20251045.

Ilse Lindenlaub, Ryungha Oh, Maria Alejandra Rodriguez, and Laura Veldkamp. Beyond exposure: Predicting AI adoption based on comparative advantage. Technical report, National Bureau of Economic Research, May 2026. URL <http://dx.doi.org/10.3386/w35271>.

Maxim Massenkoff and Peter McCrory. Labor market impacts of AI: A new measure and early evidence. Anthropic Economic Research, March 2026. URL <https://www.anthropic.com/research/labor-market-impacts>. Published March 5, 2026.

Kristina McElheran, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415, 2024. doi: <https://doi.org/10.1111/jems.12576>.

- Bruce D. Meyer and Nikolas Mittag. Misclassification in binary choice models. *Journal of Econometrics*, 200(2):295–311, 2017. doi: <https://doi.org/10.1016/j.jeconom.2017.06.012>.
- Guiseppe Pulito, Mariola Pytikova, Sarah Schroeder, and Magnus Lodefalk. Who adopts AI? Evidence on firms, technologies and workers. Technical report, IZA Discussion Paper No. 18515, 2026.
- Everett M. Rogers. *Diffusion of Innovations*. Free Press of Glencoe, New York, 1962.
- Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. Working with AI: Measuring the applicability of generative AI to occupations, July 2025. URL <https://arxiv.org/abs/2507.07935>.
- Lee C. Tucker. You’re (not) hired: Artificial intelligence and early career hiring in the quarterly workforce indicators. CES Working Paper 26-27, Center for Economic Studies, U.S. Census Bureau, April 2026.
- Erik Vasilauskas and Michael Horrigan. AI exposure and the future of work: Tasks, skill demand, and education. Fiscal policy forum policy and research paper series, The George Washington University Institute for Korean Studies, 2026. URL <https://research.upjohn.org/externalpapers/136/>.
- Michelle Yin, Hoa Vu, and Claudia Persico. How (un)stable are LLM occupational exposure scores? Evidence from multi-model replication. NBER Working Paper 35110, National Bureau of Economic Research, April 2026.

Table 1. Exposure and adoption measures

Label	Definition	Reference
Main exposure measures (occupation-level)		
drb	share of tasks completed $\geq 50\%$ faster with LLMs & software, GPT rated	Eloundou et al. (2024)
oe	observed exposure: sufficient work-related Claude traffic is observed in task	Massenkoff and McCrory (2026)
Additional exposure measures (occupation-level)		
dra	share of tasks completed $\geq 50\%$ faster with LLMs, GPT rated	Eloundou et al. (2024)
drg	share of tasks completed $\geq 50\%$ faster with LLMs & software, GPT rated	Eloundou et al. (2024)
gec	the occupation’s core tasks are replaceable by Generative AI	Eisfeldt et al. (2026)
ges	the occupation’s supplemental tasks are replaceable by Generative AI	Eisfeldt et al. (2026)
get	the occupation’s tasks are replaceable by Generative AI	Eisfeldt et al. (2026)
hra	share of tasks completed $\geq 50\%$ faster with LLMs, OpenAI human rated	Eloundou et al. (2024)
hrb	share of tasks completed $\geq 50\%$ faster with LLMs & software, OpenAI human rated	Eloundou et al. (2024)
hrg	share of tasks completed $\geq 50\%$ faster with LLMs & software, OpenAI human rated	Eloundou et al. (2024)
aioe	occupation-level measures of exposure to 10 different types of AI applications	Felten et al. (2021)
ig_aioe	image generation: links AI functions to 52 human abilities to perform jobs	Felten et al. (2023)
lm_aioe	text generation by LLMs: links AI functions to 52 human abilities to perform jobs	Felten et al. (2023)
aiapp	AI applicability score: weighted index of work activities that Copilot can assist/perform	Tomlinson et al. (2025)
Adoption measures (firm-level)		
Post-revision	indicator of AI use in any business function	Business Trends and Outlook Survey
Pre-revision	indicator of AI use in producing goods or services	Business Trends and Outlook Survey

Note: Pre-revision adoption is based on the core AI use question in the Business Trends and Outlook Survey (BTOS) for the period Sep. 2023–Oct. 2025. Post-revision adoption is based on the core AI use question in BTOS starting in Nov. 2025.

Table 2. Summary statistics

(a) Post-revision sample

	Emp-weighted			Firm-weighted		
	mean	sd	cv	mean	sd	cv
Employment	286.7	573.7	2.001	13.25	60.19	4.543
Payroll/worker	43.56	40.94	0.940	41.57	75.12	1.807
Firm age	22.14	14.16	0.640	16.54	12.74	0.770
Adoption	0.274	0.446	1.628	0.211	0.408	1.934
Exposure measures:						
drb	0.320	0.162	0.506	0.369	0.188	0.509
oe	0.132	0.121	0.917	0.161	0.153	0.950
aioe	0.473	0.237	0.501	0.532	0.268	0.504
ig_aioe	0.422	0.234	0.555	0.488	0.254	0.520
lm_aioe	0.495	0.228	0.461	0.541	0.262	0.484
aiapp	0.163	0.069	0.423	0.179	0.079	0.441
dra	0.140	0.099	0.707	0.158	0.136	0.861
drg	0.500	0.255	0.510	0.581	0.291	0.501
gec	0.183	0.128	0.699	0.209	0.149	0.713
ges	0.067	0.054	0.806	0.086	0.085	0.988
get	0.250	0.161	0.644	0.295	0.204	0.692
hra	0.143	0.091	0.636	0.160	0.119	0.744
hrb	0.287	0.143	0.498	0.321	0.173	0.539
hrg	0.432	0.208	0.481	0.482	0.247	0.512
Observations	15,500					

(b) Pre-revision sample

	Emp-weighted			Firm-weighted		
	mean	sd	cv	mean	sd	cv
Employment	293.1	898.8	3.067	12.58	59.40	4.722
Payroll/worker	44.36	48.41	1.091	41.57	83.57	2.010
Firm age	21.78	14.19	0.652	16.52	12.80	0.775
Adoption	0.090	0.285	3.167	0.069	0.253	3.667
Exposure measures:						
drb	0.317	0.164	0.517	0.366	0.190	0.519
oe	0.128	0.121	0.945	0.163	0.157	0.963
Observations	57,000					

Note: The number of observations is rounded for disclosure avoidance. “cv” denotes the coefficient of variation ($cv = sd/mean$).

Table 3. Regression estimates, robustness to controls

	Outcome: Adoption					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: GPT-4 beta (Eloundou et al., 2024)						
Exposure (drb)	0.607*** (0.0585)	0.605*** (0.0588)	0.383*** (0.0670)	0.358*** (0.0692)	0.336*** (0.0691)	0.385*** (0.0652)
Adj. R2	0.048	0.049	0.076	0.090	0.122	0.146
Within Adj. R2		0.048	0.012	0.010	0.007	0.009
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)						
Exposure (oe)	0.648*** (0.0686)	0.647*** (0.0683)	0.338*** (0.0732)	0.319*** (0.0743)	0.274*** (0.0767)	0.334*** (0.0728)
Adj. R2	0.031	0.032	0.071	0.086	0.119	0.142
Within Adj. R2		0.031	0.006	0.005	0.003	0.005
Observations	15,500	15,500	15,500	15,500	15,500	15,500
Year FEs		Yes	Yes	Yes	Yes	Yes
Industry FEs			sector	NAICS-3	NAICS-4	NAICS-4
Firm controls						Yes

Note: The table reports estimates from regressing AI adoption on AI exposure, as specified in Equation 8. The number of observations is rounded for disclosure avoidance. The sample is all single-unit firms who responded to the 2025 wave of the BTOS and one of the 2018-2020 panels of the OEWS. The outcome for each regression is an indicator for whether the firm reported using AI to perform any business function in the BTOS. The exposure variable in each panel is calculated by applying each occupation's AI exposure to the firm's occupational mix, as reported in the OEWS. Firm controls include additional fixed effects for firm age, size-class, and state. Observations in the regression are weighted so the overall sample matches the employment-weighted distribution of single-unit firms by industry and size. Stars indicate the level of significance: * 5%, ** 1%, and *** 0.1%.

Table 4. Regression estimates, heterogeneity by firm size

	Outcome: Adoption				
	(1)	(2)	(3)	(4)	(5)
Firm size class:	1 to 4	5 to 9	10 to 19	20 to 99	100+
Panel A: GPT-4 beta (Eloundou et al., 2024)					
Exposure (drb)	0.472*** (0.046)	0.512*** (0.055)	0.586*** (0.054)	0.674*** (0.051)	0.742*** (0.143)
Adj. R2	0.052	0.040	0.045	0.056	0.062
Avg Exposure	0.387	0.351	0.326	0.302	0.310
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)					
Exposure (oe)	0.499*** (0.062)	0.512*** (0.070)	0.648*** (0.081)	0.791*** (0.071)	0.728*** (0.179)
Adj. R2	0.039	0.025	0.034	0.041	0.030
Avg Exposure	0.177	0.146	0.130	0.114	0.134
Observations	2,400	3,200	3,700	5,000	1,200
Avg Adoption	0.203	0.214	0.207	0.259	0.332

Note: The table reports bivariate estimates from regressing AI adoption on AI exposure, estimated separately by the firm's size class. The sample, variable definitions, and weights match those of Table 3, Column 1. The number of observations is rounded for disclosure avoidance. Stars indicate the level of significance: * 5%, ** 1%, and *** 0.1%.

Table 5. Regression estimates, heterogeneity by firm sector

Sector	Outcome: Adoption														
	(1) 72	(2) 48–49	(3) 81	(4) 23	(5) 42	(6) 62	(7) 44–45	(8) 31–33	(9) 71	(10) 56	(11) 61	(12) 53	(13) 52	(14) 54	(15) 51
Panel A: GPT-4 beta (Eloundou et al., 2024)															
Exposure (drb)	-0.205 (0.433)	0.771 (0.354)	0.441 (0.239)	0.660 (0.184)	0.637 (0.170)	0.414 (0.181)	0.787 (0.276)	0.931 (0.214)	0.580 (0.381)	0.204 (0.257)	-0.112 (0.660)	0.758 (0.382)	1.074 (0.575)	0.697 (0.143)	2.144 (0.431)
Adj. R2	-0.001	0.044	0.015	0.023	0.035	0.010	0.031	0.039	0.015	0.006	-0.004	0.030	0.014	0.029	0.163
Avg Exposure	0.196	0.307	0.255	0.214	0.380	0.284	0.329	0.270	0.290	0.273	0.410	0.367	0.571	0.520	0.584
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)															
Exposure (oe)	0.048 (0.417)	1.017 (0.429)	0.877 (0.325)	0.690 (0.316)	0.628 (0.160)	-0.195 (0.204)	0.615 (0.267)	0.787 (0.284)	0.680 (0.518)	0.370 (0.328)	-0.067 (0.440)	0.429 (0.481)	-0.160 (0.715)	0.832 (0.175)	1.427 (0.421)
Adj. R2	-0.003	0.049	0.027	0.006	0.030	0.001	0.016	0.017	0.010	0.011	-0.004	0.002	-0.003	0.038	0.070
Avg Exposure	0.041	0.065	0.078	0.057	0.194	0.143	0.164	0.090	0.106	0.117	0.191	0.159	0.312	0.221	0.264
Observations	400	350	550	1,300	850	1,400	1,100	1,800	350	550	250	250	250	1,200	200
Avg Adoption	0.140	0.150	0.154	0.206	0.235	0.245	0.250	0.255	0.346	0.347	0.371	0.382	0.403	0.473	0.582

Note: The table reports bivariate estimates from regressing AI adoption on AI exposure, estimated separately by the firm's sector. The sample is all single-unit firms who responded to the 2025 wave of the BTOS and one of the 2018-2020 panels of the OEWS. We additionally restrict the sample to firms with at least 5 employees and whose employment in the LBD and the OEWS differ by no more than 50%. The variable definitions, and weights match those of Table 3, Column 1. The number of observations is rounded for disclosure avoidance. Columns are ordered by average adoption from lowest to highest.

NAICS sectors: (72) Accommodation and Food Services; (48–49) Transportation and Warehousing; (81) Other Services (except Public Administration); (23) Construction; (42) Wholesale Trade; (62) Health Care and Social Assistance; (44–45) Retail Trade; (31–33) Manufacturing; (71) Arts, Entertainment, and Recreation; (56) Administrative and Support and Waste Management and Remediation Services; (61) Educational Services; (53) Real Estate and Rental and Leasing; (52) Finance and Insurance; (54) Professional, Scientific, and Technical Services; (51) Information.

Figure 1. Conceptual relationship between firm-level AI exposure, adoption, and outcomes






Figure 2. Regression estimates, all exposure measures

(a) Without controls




(b) With year and NAICS-3 fixed-effects




Note: The figure plots coefficient estimates and 95% confidence intervals from regressing AI adoption on AI exposure. Panel A reports coefficients from a bivariate regression without any controls, and Panel B reports coefficients with year and NAICS-3 fixed effects. The sample and weights match those of Table 3. The adjusted R^2 for each regression is reported to the right.

Figure 3. Regression estimates, all exposure measures (standardized)

(a) Without controls




(b) With year and NAICS-3 fixed-effects



Note: The figure plots normalized coefficient estimates from regressing AI adoption on AI exposure. The sample, weights, and specification of the figure are the same as in Figure 2, but all coefficients are scaled by the standard deviation of their respective exposure measure.

Figure 4. Predicted AI adoption by sector



Note: The figure reports average versus predicted rates of AI adoption by sector. The sample is all single-unit firms who responded to the 2025 wave of the BTOS and one of the 2018-2020 panels of the OEWS. We additionally restrict the sample to firms with at least 5 employees and whose employment in the LBD and the OEWS differ by no more than 50%. The actual adoption rate is calculated directly from this sample. The predicted adoption rate is calculated by mapping the same sample's average exposure, also reported in Table 5, to adoption using the economy-wide bivariate relationship between adoption and exposure. Sectors are ordered by actual adoption rate from lowest to highest.

NAICS sectors: (72) Accommodation and Food Services; (48-49) Transportation and Warehousing; (81) Other Services (except Public Administration); (23) Construction; (42) Wholesale Trade; (62) Health Care and Social Assistance; (44-45) Retail Trade; (31-33) Manufacturing; (71) Arts, Entertainment, and Recreation; (56) Administrative and Support and Waste Management and Remediation Services; (61) Educational Services; (53) Real Estate and Rental and Leasing; (52) Finance and Insurance; (54) Professional, Scientific, and Technical Services; (51) Information.

Appendix: Additional tables and figures

Table A1. Pairwise correlations between exposure measures

	aioe	lm_aioe	ig_aioe	get	gec	ges	dra	drb	drg	hra	hrb	hrg	aiapp
aioe													
lm_aioe	0.98												
ig_aioe	0.85	0.74											
get	0.78	0.76	0.64										
gec	0.76	0.74	0.66	0.94									
ges	0.39	0.39	0.26	0.59	0.28								
dra	0.40	0.40	0.28	0.64	0.58	0.43							
drb	0.84	0.82	0.72	0.87	0.83	0.49	0.72						
drg	0.89	0.86	0.79	0.82	0.80	0.43	0.46	0.95					
hra	0.68	0.71	0.49	0.68	0.64	0.38	0.45	0.70	0.69				
hrb	0.85	0.85	0.70	0.83	0.81	0.42	0.50	0.88	0.90	0.88			
hrg	0.86	0.85	0.74	0.83	0.81	0.41	0.48	0.89	0.92	0.74	0.97		
aiapp	0.69	0.72	0.50	0.67	0.65	0.35	0.49	0.71	0.68	0.69	0.75	0.72	
oe	0.54	0.54	0.43	0.65	0.62	0.36	0.49	0.63	0.58	0.49	0.59	0.59	0.53

Note: This table reports pairwise correlations between various AI exposure measures computed using the OEWS (2019). A darker shading indicates a higher correlation.

Table A2. Regression estimates, high-quality matches

	Outcome: Adoption					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: GPT-4 beta (Eloundou et al., 2024)						
Exposure (drb)	0.637*** (0.0760)	0.633*** (0.0758)	0.460*** (0.0879)	0.438*** (0.0917)	0.443*** (0.0770)	0.529*** (0.0692)
Adj. R2	0.045	0.047	0.064	0.075	0.110	0.145
Within Adj. R2		0.045	0.014	0.012	0.008	0.012
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)						
Exposure (oe)	0.658*** (0.0933)	0.654*** (0.0925)	0.366*** (0.0979)	0.340*** (0.100)	0.255** (0.0932)	0.319*** (0.0856)
Adj. R2	0.027	0.029	0.057	0.068	0.105	0.138
Within Adj. R2		0.027	0.006	0.005	0.002	0.003
Observations	11,000	11,000	11,000	11,000	11,000	11,000
Year FEs		Yes	Yes	Yes	Yes	Yes
Industry FEs			sector	NAICS-3	NAICS-4	NAICS-4
Firm controls						Yes

Note: The table reports estimates from regressing AI adoption on AI exposure, as specified in Equation 8. The sample is all single-unit firms who responded to the 2025 wave of the BTOS and one of the 2018-2020 panels of the OEWS. We additionally restrict the sample to firms with at least 5 employees and whose employment in the LBD and the OEWS differ by no more than 50%. The variable definitions and weights match those of Table 3. Stars indicate the level of significance: * 5%, ** 1%, and *** 0.1%.

Table A3. Regression estimates, firm-weighted

	Outcome: Adoption					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: GPT-4 beta (Eloundou et al., 2024)						
Exposure (drb)	0.445*** (0.0363)	0.446*** (0.0364)	0.244*** (0.0445)	0.215*** (0.0463)	0.210*** (0.0494)	0.256*** (0.0473)
Adj. R2	0.042	0.042	0.067	0.076	0.092	0.120
Within Adj. R2		0.042	0.008	0.006	0.005	0.007
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)						
Exposure (oe)	0.482*** (0.0488)	0.482*** (0.0488)	0.326*** (0.0544)	0.299*** (0.0554)	0.258*** (0.0568)	0.298*** (0.0541)
Adj. R2	0.033	0.033	0.070	0.079	0.093	0.121
Within Adj. R2		0.033	0.012	0.010	0.006	0.008
Observations	15,500	15,500	15,500	15,500	15,500	15,500
Year FEs		Yes	Yes	Yes	Yes	Yes
Industry FEs			sector	NAICS-3	NAICS-4	NAICS-4
Firm controls						Yes


Note: The table reports estimates from regressing AI adoption on AI exposure, as specified in Equation 8. The sample and variable definitions match those of Table 3. Observations in the regression are weighted so the overall sample matches the firm-weighted distribution of single-unit firms by industry and size. Stars indicate the level of significance: * 5%, ** 1%, and *** 0.1%.

Table A4. Regression estimates, pre-revision adoption

	Outcome: Adoption					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: GPT-4 beta (Eloundou et al., 2024)						
Exposure (drb)	0.326*** (0.0309)	0.324*** (0.0310)	0.240*** (0.0242)	0.218*** (0.0283)	0.154*** (0.0285)	0.187*** (0.0248)
Adj. R2	0.035	0.041	0.053	0.063	0.081	0.104
Within Adj. R2		0.035	0.012	0.009	0.004	0.005
Panel B: Observed Exposure (Massenkoff and McCrory, 2026)						
Exposure (oe)	0.356*** (0.0343)	0.357*** (0.0338)	0.246*** (0.0266)	0.222*** (0.0301)	0.155*** (0.0331)	0.183*** (0.0293)
Adj. R2	0.023	0.030	0.050	0.061	0.081	0.103
Within Adj. R2		0.023	0.008	0.006	0.003	0.004
Observations	57,000	57,000	57,000	57,000	57,000	57,000
Year FEs		Yes	Yes	Yes	Yes	Yes
Industry FEs			sector	NAICS-3	NAICS-4	NAICS-4
Firm controls						Yes

Note: The table reports estimates from regressing AI adoption on AI exposure, as specified in Equation 8. The sample is all single-unit firms that responded to one of the 2023-2025 waves of the BTOS and one of the 2018-2020 panels of the OEWS. The outcome for each regression is an indicator for whether the firm reported using AI to produce goods or services in the BTOS. The exposure variable in each panel is calculated by applying each occupation's AI exposure to the firm's occupational mix, as reported in the OEWS. Firm controls include additional fixed effects for firm age, size-class, and state. Observations in the regression are weighted so the overall sample matches the employment-weighted distribution of single-unit firms by industry and size. Stars indicate the level of significance: * 5%, ** 1%, and *** 0.1%.


Figure A1. Regression estimates, all exposure measures, pre-revision adoption




Note: The figure plots coefficient estimates and 95% confidence intervals from regressing AI adoption on AI exposure. Panel A reports coefficients from a bivariate regression without any controls, and Panel B reports coefficients with year and NAICS-3 fixed effects. The sample and weights match those of Table A4. The adjusted R^2 for each regression is reported to the right.

Figure A2. Regression estimates, heterogeneity by sector

(a) GPT-4 beta (Eloundou et al., 2024)



(b) Observed Exposure (Massenkoff and McCrory, 2026)



Note: The figure plots the coefficient estimates and 95% confidence intervals for the regressions from Table 5. The adjusted R^2 for each regression is reported to the right. For definitions of sector codes, see the notes for Table 5.