

AI Exposure and Organizational Structure

Guohou Shan

D'Amore-McKim School of Business, Northeastern University

360 Huntington Ave, Boston, MA 02115,

g.shan@northeastern.edu

Feng Zhu

Harvard Business School

Massachusetts Hall, Cambridge, MA 02138

fzhu@hbs.edu

AI Exposure and Organizational Structure

Abstract

The rapid diffusion of artificial intelligence (AI) has intensified debate about whether AI will flatten organizational hierarchies by substituting for managerial functions or instead expand them by generating new coordination and oversight demands. Despite strong claims from executives and consultants, evidence on how AI exposure reshapes core organizational structures remains limited. We address this gap by examining whether and how AI exposure affects two fundamental dimensions of organizational design: hierarchy depth and managerial span of control. We exploit staggered exposure to AI, defined as the presence of employees with AI-related skills, across 7,681 U.S. firms between 1990 and 2020 and implement a difference-in-differences research design. We find that AI exposure is associated with systematic restructuring of organizations. Firms exposed to AI reduce the number of hierarchical ranks while simultaneously increasing managerial spans of control, consistent with a flattening of organizational hierarchies. These effects are economically meaningful and heterogeneous: they are stronger in technical firms, in industries with higher automation intensity, and in organizations with greater technical expertise among managers and employees. Our findings contribute to the literature on organizational design and technology by moving beyond individual-level productivity effects to show how AI reshapes the internal structure of firms. For managers and policymakers, the results provide evidence-based guidance on how AI adoption interacts with hierarchy and managerial scope, thereby helping organizations redesign structures in ways that are both technologically enabled and organizationally sustainable.

Keywords: AI exposure, organizational structure, hierarchical ranks, span of control, difference-in-differences

1. Introduction

The rapid diffusion of AI across organizations has triggered intense debate among executives, consultants, and policymakers about how organizations should redesign their internal structures. Industry leaders increasingly argue that AI will “flatten” organizations by reducing managerial layers while enabling leaders to oversee larger teams through algorithmic decision support, real-time analytics, and automation of coordination tasks (Anthony et al. 2023; Nolan 2025). In contrast, others caution that AI may introduce new forms of complexity, requiring additional oversight, specialized roles, and governance mechanisms that could reinforce or even expand hierarchical structures (Raisch and Krakowski 2021; Trustible 2024; Rai et al. 2026). Despite widespread adoption and strategic urgency, there remains little systematic empirical evidence on whether AI exposure actually reshapes organizations’ hierarchies and, if so, how.

This debate reflects a deeper theoretical tension in the organizational design literature. Classic work emphasizes that hierarchy and span of control are shaped by information-processing constraints and coordination costs (Simon 1960; Galbraith 1974). Digital technologies have historically altered these constraints, often enabling flatter structures by improving information flow and monitoring (e.g., Bloom et al. 2014). AI represents a qualitatively distinct technological shift because it augments cognitive and analytical tasks traditionally performed by managers (Berente et al. 2021), raising the possibility that it substitutes for middle management while allowing remaining managers to supervise more employees. When organizations are exposed to AI, for example, by hiring AI talent, they begin to leverage it to improve information processing (Babina et al., 2024). Yet existing empirical research on AI has largely focused on individual productivity, task performance such as innovation, or skill demand (Lou et al. 2021; Wang et al. 2023; Brynjolfsson et al. 2025; Lou et al. 2025; Wu et al. 2025; Tambe 2026; Wang and Wu 2026), leaving the organizational-structure implications underexplored.

In this paper, we examine whether AI exposure leads to structural flattening in organizations. Consulting reports and executive commentaries frequently predict that AI-enabled decision support will eliminate intermediate managerial layers, reducing the number of hierarchical ranks while increasing spans

of control (McKinsey Global Institute 2023). However, these claims are largely speculative and rest on case anecdotes rather than large-scale evidence. The academic literature has yet to provide systematic causal estimates of whether AI exposure is associated with measurable changes in core organizational design outcomes, such as hierarchy depth and managerial span.

This paper addresses this gap by providing the first large-scale longitudinal evidence on whether AI exposure is associated with changes in organizational ranks and spans of control. From a managerial perspective, this question is critical: decisions about AI adoption are increasingly intertwined with workforce restructuring, managerial layoffs, and redesign of reporting relationships. Clear evidence on whether AI exposure systematically flattens or reshapes hierarchies can help executives move beyond conjecture when planning organizational redesigns and managing employee expectations during AI-driven transformations.

Beyond whether AI matters, industry debates increasingly emphasize how its effects unfold and under what conditions. Executives report heterogeneous outcomes: some firms successfully expand managerial spans using AI tools, while others struggle with overload, resistance, or coordination failures (Davenport et al. 2020; Raisch and Krakowski 2021). This heterogeneity suggests that AI's organizational impact is not uniform but may depend on complementary factors such as technical human capital, industry-level automation intensity, and industry-level technical depth.

Prior studies often treat technology adoption as homogeneous, overlooking how workforce capabilities and organizational context shape its consequences (Bloom et al. 2014; Brynjolfsson et al. 2025; Cheng et al. 2025). Understanding the boundary conditions can inform more effective implementation strategies. Rather than assuming AI will automatically flatten organizations, managers can better anticipate where complementary investments, such as technical training or process redesign, are necessary to realize AI's promised organizational benefits.

To answer these questions, we assemble a large-scale longitudinal dataset covering 7,681 U.S. firms between 1990 and 2020 by matching detailed workforce and organizational data from Revelio Labs with firm-level information used in prior organizational research (e.g., Li et al. 2022; Liang et al. 2024, 2025;

Marchetti and Puranam 2025). We exploit staggered adoption of AI across firms and implement a difference-in-differences (DID) research design, defining AI exposure as the presence of at least one employee with AI-related skills, following Babina et al. (2024). This approach allows us to identify changes in organizational structure associated with AI exposure while controlling for time-invariant firm characteristics and common temporal shocks.

Our findings reveal a systematic restructuring of organizations following AI exposure. Firms exposed to AI reduce the number of hierarchical ranks while simultaneously increasing spans of control, measured by the number of direct reports. These effects are amplified in organizations with stronger technical backgrounds among managers and employees, and in highly technical or highly automated industries.

By examining both whether and how AI exposure reshapes organizational hierarchies, this study directly informs ongoing industry debates about the “end of middle management” and the future of organizational design in the AI era (e.g., Gavett and Sawhney 2025; Tremblay 2025). Conceptually, we extend the AI and information systems (IS) literature beyond productivity and task outcomes to the foundational question of organizational structure. For practitioners, our findings provide evidence-based guidance on how AI adoption interacts with hierarchy, managerial scope, and workforce composition, helping organizations redesign structures that are not only technologically advanced but also organizationally sustainable.

2. Related Literature and Theoretical Foundation

In this section, we discuss three related streams of research: 1) AI and organizational design, 2) hierarchy, span of control, and information processing, and 3) AI as managerial augmentation.

2.1 AI and Organizational Design

A growing body of research in information systems and economics examines the consequences of AI and advanced digital technologies for work and organizational outcomes. Much of this literature focuses on individual productivity, task performance, and skill demand, showing that AI can augment human labor,

reallocate tasks, and reshape occupational structures (Acemoglu and Restrepo 2018; Cheng et al. 2021; Tschang and Almirall 2021; Xue et al. 2022; Brynjolfsson et al. 2025; Qiao et al. 2026). These studies provide important insights into how AI affects *what* work is done and *who* does it, but they largely abstract away from how work is *organized* within firms.

Earlier work on information technology (IT) and organizational design suggests that digital technologies can reduce coordination and monitoring costs, thereby enabling flatter organizational structures (Bloom et al. 2014). However, AI differs from earlier generations of IT in that it directly supports cognitive, analytical, and decision-making activities traditionally performed by managers (Shrestha et al. 2019). As a result, AI has the potential not only to improve information flow but also to substitute for managerial functions such as evaluation, planning, and supervision. Despite this potential, empirical evidence on how AI affects core structural features, such as the number of hierarchical ranks and spans of control, remains sparse.

Recent conceptual work highlights this gap. Scholars argue that AI may simultaneously automate managerial tasks and create new forms of organizational complexity, leading to ambiguous effects on hierarchy (Raisch and Krakowski 2021). Yet these arguments have not been tested systematically using large-scale firm-level data. Consequently, while industry discussions frequently assert that AI will “flatten” organizations, the academic literature has not yet established whether AI exposure is associated with observable changes in organizational hierarchies.

2.2 Hierarchy, Span of Control, and Information Processing

Classic theories of organizational design view hierarchy and span of control as mechanisms for managing information-processing demands (Simon 1960; Galbraith 1974). Hierarchical layers help aggregate information and resolve coordination problems, while spans of control reflect the extent to which managers can effectively oversee subordinates (Keren and Levhari 1979; Mihm et al. 2010; Gómez et al. 2016). When information-processing capacity is limited, organizations tend to adopt taller structures with narrower spans of control (Wang and Li 2025). Conversely, improvements in information-processing technologies can enable flatter structures and wider managerial spans.

Empirical research supports these ideas. Studies show that improvements in communication and monitoring technologies allow managers to supervise more employees and reduce the need for intermediate layers (Bloom et al. 2014). However, this literature largely predates recent advances in AI and focuses on traditional IT rather than algorithmic decision-support systems. AI may alter information-processing constraints more profoundly by automating analysis, generating recommendations, and standardizing decisions across organizational levels (Jarrahi 2018; Shrestha et al. 2019).

At the same time, organizational design research emphasizes that structural changes are contingent on organizational context. Larger firms, for example, face greater coordination challenges and may benefit more from technologies that expand managerial capacity (Child 1987). Similarly, industries characterized by high levels of automation or technical complexity may be better positioned to integrate AI into managerial workflows (Spring 2022). These insights suggest that AI's effects on hierarchy and span of control are unlikely to be uniform across firms, underscoring the importance of examining moderating factors such as firm size, workforce technical background, and industry characteristics.

2.3 AI as Managerial Augmentation

Building on the information-processing view of organizations, we conceptualize AI as a form of managerial augmentation technology. AI systems can process large volumes of data, provide real-time performance feedback, and support standardized decision-making, thereby expanding managers' effective information-processing capacity. From this perspective, AI exposure should reduce the need for additional hierarchical layers by enabling managers to oversee more direct reports without proportional increases in cognitive burden.

However, this augmentation effect depends on complementary organizational resources. Managers and employees with stronger technical backgrounds may be better able to interpret and trust AI-generated insights (Cohen and Mahabadi 2022), amplifying AI's impact on spans of control. Similarly, larger organizations may realize greater structural benefits from AI because scale increases the returns to coordination efficiencies (Fountain et al. 2019). In contrast, in contexts lacking technical expertise or facing institutional constraints, AI may fail to substitute for managerial layers, limiting its structural impact.

The integration of insights from the AI literature with classic organizational design theory provide theoretical foundation for examining how AI exposure affects hierarchy and span of control. It also provides a foundation for examining heterogeneity in AI’s organizational consequences, allowing us to move beyond the question of whether AI matters to understand how and under what conditions it reshapes organizational structure.

3. Empirical Context and Data

3.1 Research Context

We examine our research questions in the context of the U.S. labor market, leveraging large-scale employer–employee data from Revelio Labs. The United States provides a particularly suitable empirical setting for studying the organizational consequences of AI adoption because U.S. firms have been early and intensive adopters of AI-related technologies across a wide range of industries (McElheran et al. 2024). Moreover, the availability of detailed workforce data allows for systematic measurement of both AI exposure and organizational structure over time.

3.2 Data Sources and Sample Construction

To address our research questions, we rely on large-scale employer–employee data from the U.S. labor market, drawing primarily on Revelio Labs (henceforth, Revelio). Revelio is a workforce intelligence company that has collected and harmonized information from over 380 million online employee profiles and resumes since the early 2000s. Revelio data have become increasingly popular in academic research (Li et al. 2022; Liang et al. 2024, 2025; Marchetti and Puranam 2025) and business research (Sull et al. 2019, 2022) for enhancing our understanding of how firms organize, grow, and adapt to technological change.

Revelio provides individual-level longitudinal information on employees’ historical job positions, job titles, education, skills, and career trajectories. These data allow us to reconstruct firms’ internal organizational structures and track changes in workforce composition over time. Importantly for our setting,

Revelio predicts employees' seniority levels based on a combination of age, current job title, previous seniority, career history, and firm- and industry-level characteristics. Seniority is coded on a seven-level scale, where 1 represents the most junior roles and 7 the most senior. More specifically, according to Revelio dictionary (Revelio Labs 2024), level 1 is the entry level (Ex. Accounting Intern, Software Engineer Trainee, Paralegal), level 2 is the junior level (Ex. Account Receivable Bookkeeper, Junior Software QA Engineer, Legal Adviser), level 3 is the associate level (Ex. Senior Tax Accountant; Lead Electrical Engineer; Attorney), level 4 is the manager level (Ex. Account Manager; Superintendent Engineer; Lead Lawyer), level 5 is the director level (Ex. Chief of Accountants; VP Network Engineering; Head of Legal), level 6 is the executive level (Ex. Managing Director, Treasury; Director of Engineering, Backend Systems; Attorney, Partner), and level 7 is the senior executive level (Ex. CFO; COO; CEO). We leverage this information to infer reporting relationships within firms.

Our empirical sample consists of 7,681 U.S. firms observed annually between 1990 and 2020. We restrict the sample to firms for which Revelio provides sufficient workforce coverage to reliably infer organizational structure. Following prior work that uses job titles and seniority information to estimate organizational hierarchies (Rajan and Wulf 2006; Colombo and Delmastro 2008), we aggregate individual-level data to the firm-year level. The resulting panel structure allows us to examine how firms' internal hierarchies evolve over time, particularly around the onset of AI exposure.

3.3 AI Exposure Identification

We identify firms' exposure to AI following the approach developed by Babina et al. (2024). Specifically, we create a dummy variable, $AIExposure_{it}$, to indicate whether firm i employs at least one worker whose skill set includes AI-related capabilities in year t . These skills are identified using detailed skill descriptors contained in individual professional profiles, covering areas such as machine learning, artificial intelligence, deep learning, and related computational techniques.

This measure captures the point at which AI-related human capital becomes embedded within the firm's workforce, reflecting meaningful organizational exposure to AI rather than indirect use through vendors or external consultants. The AI exposure indicator is time-variant and allows us to exploit staggered adoption

of AI across firms. Once a firm first becomes AI-exposed, it remains classified as exposed in subsequent years, consistent with the idea that AI-related capabilities represent a persistent organizational investment. This staggered timing serves as the basis for our difference-in-differences empirical strategy.

3.4 Measures of Organizational Structure

Our primary dependent variables capture two core dimensions of organizational structure: hierarchical ranks and span of control. Following Rothaermel (2018) and Vergne (2020), we conceptualize organizational hierarchy as the chain of command within the firm—that is, the extent to which authority and decision-making are distributed across distinct vertical levels.

Hierarchical ranks, $Rank_{it}$, measure the number of distinct organizational layers within firm i in year t . Using Revelio’s predicted seniority classifications, we identify the number of vertical levels separating the most junior employees from the most senior managers. Revelio classifies an organization’s hierarchical ranks into seven seniority levels covering all organizational hierarchy (Revelio Labs 2024), which has also been adopted by other papers, such as Huang et al. (2025). Firms with more hierarchical ranks exhibit taller organizational structures characterized by greater segmentation of authority and decision-making across levels.

Span of control, $SpanControl_{it}$, captures the breadth of managerial oversight and is measured as the average number of direct reports per manager within a firm i in year t . Managers are identified based on seniority levels and job titles, following established approaches in the organizational design literature. Wider spans of control indicate flatter structures in which managers oversee larger teams, whereas narrower spans suggest more intensive supervision and a greater reliance on intermediate managerial layers.

Together, these measures provide a comprehensive view of how authority and coordination are organized within firms. By combining them with a time-variant measure of AI exposure, we are able to examine whether and how the introduction of AI-related human capital is associated with systematic changes in firms’ internal hierarchies.

3.5 Panel Structure

We organize the data at the firm-year level, resulting in an unbalanced panel spanning three decades. This

structure allows us to observe firms both before and after AI exposure and to control for time-invariant firm characteristics and common temporal shocks. We also control for firm size, $Size_{it}$, measured as the number of employees of firm i in year t . It is log-transformed to avoid data-skewed distribution (Lu and Lu 2022). It is a time-variant variable and a determinant of an organizational hierarchy, according to classic literature (e.g., Marsh and Mannari 1981). Summary statistics of these variables are reported in Table 1.

From Table 1, the AI exposure is 0.095, indicating that 9.5% of observations involve AI exposure. The mean values of rank and span of control are 6.333 and 3.33, and their log transformations are 1.967 and 1.370, respectively. In addition, the log-transformed average firm size is 3.649, which corresponds to the average log-transformed number of employees per firm per year.

The unit of analysis in this study is a firm-year. In other words, all our variables are measured at the yearly level for each firm in our sample. We used log transformation for the dependent variables and ran Poisson models because of the count data and got consistent results. For public firms, we controlled for financial performance because organizational structure and financial performance are highly correlated (Guadalupe et al. 2014), and we obtained consistent results. In addition, given the potential latency of the AI exposure effect, we lagged the dependent variable and obtained consistent results. More details are presented in the robustness check section.

4. Empirical Results

4.1 Main Results

In this study, our main dependent variables are hierarchical rank and span of control, measured at the firm level annually. To investigate the main impact, we organize a DID model with staggered adoption of the firm's AI exposure. We also use a two-way fixed effects (TWFE) panel model to estimate the main effect (Wei and Lin 2017; Wang et al. 2022; Pu et al. 2024). The model specifications are shown in equation (1):

$$\{Rank_{it}, SpanControl_{it}\} = \alpha_t + \gamma_i + \beta_1 AIExposure_{it} + Control_{it} + \varepsilon_{it}, \quad (1)$$

where $AIExposure_{it}$ is a dummy variable indicating whether firm i has hired employees whose skill set

includes AI-related capabilities until year t . $Control_{it}$ captures time-variant control variables such as the firm size as measured by the number of employees in a firm in a particular year. In addition, we incorporated the α_t and γ_i , capturing the time-level and firm-level fixed effects, respectively. The TWFE estimation results on the matches sample are presented in Table 2.

As shown in Table 2, we found that exposure to AI significantly affects firms' organizational structure. Specifically, we observe that exposure to AI decreases the number of ranks by 0.203 and increases the span of control by 0.0278, corresponding to a 3.21% ($=0.203/6.333$) decrease in the number of ranks and a 0.83% increase in the span of control, on average.

These empirical findings support the idea that when exposed to AI, firms typically leverage its capabilities to improve their information processing efficiency, leading organizations to become flatter. Additionally, each layer of employees empowered by AI can manage more direct reports, thereby increasing their span of control. These findings align with recent industrial reports that AI will “flatten” organizations by reducing managerial layers while enabling leaders to oversee larger teams through algorithmic decision support, real-time analytics, and the automation of coordination tasks (Anthony et al. 2023; Nolan 2025; Wang et al. 2025).

4.2 Robustness Checks

4.2.1 Parallel Trend Investigation

To ensure the validity of our DID estimate, we examine the parallel trend assumption. We employ leads and lags models, as suggested by Alyakoob and Rahman (2022) and Liu et al. (2024). Liu et al. (2024) caution that the coefficients and p -values of the treatment indicator and dummy variables in traditional leads and lags models may be biased due to their dependence on the selected baseline category. To overcome this limitation, we follow their approach and plot the average differences between the outcomes and the predicted counterfactual for each firm i at period t (relative to the treatment) in the treatment groups, utilizing the interactive fixed effects counterfactual (IFEct) estimator. These plots are presented in Figure 1 for the firm-level dependent variables. By examining the average pre-treatment residuals, we can

determine whether they converge to zero, which indicates the presence of a parallel trend. Our results show that during the pre-treatment periods, the average pre-treatment residuals do not deviate significantly from zero. This suggests that our DID model satisfies the parallel trend assumption, thereby strengthening the confidence in our findings.

4.2.2 Instrumental Variable Approach

Aside from using a counterfactual estimator to address time-variant confounders, we adopt an instrumental variable (IV) approach to further address endogeneity arising from unobserved selection into firms' decisions to hire employees with AI skills. Specifically, we follow prior literature and use the two-stage least squares (2SLS) method to estimate the regression model.

In our study, we introduce a relevant IV, firms' exposure to the supply of AI talent from U.S. universities (*AITalents*), constructed by Babina et al. (2024). The *AITalents* help isolate variation in firms' decisions to hire employees with AI skills that stems from the supply of AI labor, mitigating potential bias from demand shocks that drive both firms' decisions and their organizational structure changes. The limited availability of AI-skilled labor represents a central constraint on firms' adoption of AI technologies. Consequently, variation in access to AI talent is highly informative about firms' decisions to hire AI workers and initiate AI adoption. In this context, firms' preexisting connections to universities with strong AI research programs provide a plausibly exogenous source of variation in the supply of AI talent available to firms during the surge of commercial interest in AI in the 2010s, a period that followed decades of academic research in the field (Babina et al. 2023). This IV is plausibly exogenous because the supply of AI talent cannot affect firms' organizational structure unless it affects their decision to hire AI workers. Thus, both the relevance and exclusion conditions are met for the usage of the IV.

In the first stage of estimating the model and main effects, we run the analysis of *AIExposure* on the IV and control variables. In the second stage, we use *AIExposure (Instrumented)*, the fitted value of *AIExposure* from the first stage, as the main independent variable and use organizational structure as the dependent variable. The estimation results are shown in Table 3, reaffirming our main results.

We also conduct a weak instrument test. The Cragg-Donald Wald F-statistic is 1422.85, which exceeds the critical value of 16.38 (based on two-stage least-squares with one instrument, significance level at 0.05; see Stock and Yogo (2005)). Thus, our instrument, *AITalents* is not a weak instrument.

In addition, we found the significant increases in the coefficients in the 2SLS regressions, which suggest the magnitude of the effects we observe is in reality larger than suggested by the OLS results. The 2SLS estimates reflect a local average treatment effect (Imbens and Angrist 1994; Angrist and Pischke 2009), since not all firms are exposed to the instrument, and only those that hire AI talent from US universities are plausibly induced to adopt AI themselves. Therefore, in the paper, we follow Imbens and Angrist (1994) and Angrist and Pischke (2009) and restrict our discussion to the sign of these coefficients.

4.2.3 Falsification Test

One potential concern about the significant effects in our DID estimation may be false significance due to spurious relationships in our dependent variables. Following prior literature (e.g., Burtch et al. 2018), we addressed this concern by conducting a falsification test in which we randomly implemented the AI exposure to each firm in a given year (e.g., by randomly generating and assigning the treatment). Using the pseudo-treatments for each user as a placebo, we conducted regressions (equation (1)) to obtain the treatment estimates and replicated this falsification test 1,000 times. This test allows us to check that the effect we found is due to AI exposure, not to variable correlations or outliers. Our results, shown in Table 4, reveal that the mean of the random treatments is close to zero and significantly different from our estimated effect, suggesting that the impacts of our treatments are not driven by spurious relationships.

4.2.4 IFect Estimator

Using TWFE estimators, we controlled the firm and time fixed effects (i.e., time-invariant confounders). However, there are possibly unobserved variables that vary over time and are correlated with both treatment and outcomes (e.g., unobserved time-variant confounders affect firms' decision to hire employees who have an AI skill set). A group of methods has been proposed to address such potential issues (Cengiz et al. 2019; Goodman-Bacon 2021; Baker et al. 2022; De Chaisemartin and D'Haultfoeuille 2024; Liu et al. 2024). They are generally classified into three categories: the imputation estimator (e.g., Liu et al. 2024), the Cohort-

Specific Average Treatment Effects on the Treated (CATT) (e.g., De Chaisemartin and D'Haultfoeuille 2024), and the stacked regression estimator (e.g., Cengiz et al. 2019). In our research, we primarily employ the imputation estimator because it facilitates the examination of parallel trends, an important consideration in our study (Liu et al. 2024). For the imputation estimator, we follow Liu et al. (2024) and construct the interactive fixed effects counterfactual (IFect) estimator.

The imputation estimator we employed possesses the capability to address a certain degree of unobserved time-variant confounders, which is achieved through the utilization of a factor-augmented model to estimate the counterfactual potential outcomes for the untreated units. As outlined by Liu et al. (2024), this modeling approach has the potential to mitigate the impact of unobserved time-variant confounders on the differences observed between treated and untreated units. Our estimation results are shown in Table 5. The results show that the IFect estimates are consistent with the TWFE estimator, confirming that AI exposure can significantly decrease the number of ranks and increase the span of control.

4.2.5 Alternative Measurements and Lagged Effects

We conducted robustness checks on alternative measurements of the independent and dependent variables. First, we treated firms as exposed to AI when they began employing AI workers, without accounting for the number of AI workers they hired. To test whether our results are sensitive to the number of AI workers hired, we conduct a robustness check by using the number of AI workers hired each year, *AIExposure_2_{it}*, as the measurement of the AI exposure. Our results, shown in Table 6, reveal consistent findings.

Second, *Rank* and *SpanControl* may have skewed distribution (especially for the span of control, shown in Table 1). We conducted robustness checks by using the log transformation for the dependent variables (i.e., natural log). The results presented in Panel A of Table 7 reveal consistent findings. In addition, given that *Rank* and *SpanControl* are count data. We also ran Poisson models, getting consistent findings, as shown in Panel B of Table 7.

Third, given that the effect of AI exposure may have some latency, we also lag the dependent variables by one year, two years, and three years and found that the AI exposure impacts are consistent, as shown in Table 8.

Fourth, since both company size and span of control involve the number of employees, we also tested specifications that include lagged company size as a control, as well as specifications that exclude it. Our results, shown in Table 9, reveal consistent findings.

4.2.6 Time-Variant Controls

In addition to using firm size as a time-variant control, we include other time-variant controls. Literature suggests that firms' financial performance correlates with their organizational structure (e.g., Guadalupe et al. 2014). Therefore, to mitigate potential impacts of financial performance and confirm the robustness of our main findings, we include some time-variant financial controls. Following the literature (e.g., Agarwal et al. 2021), we include a set of financial performance controls for publicly traded firms from CompuStat, including advertising intensity (*Adv_sale*), measured as advertising expenses divided by sales; capitalization ratio (*Capital_ratio*), computed as total long-term debt divided by the sum of long-term debt, common/ordinary equity, and preferred stock; cash ratio (*Cash_ratio*), measured as cash and short-term investments divided by current liabilities; return on assets (*ROA*), defined as operating income before depreciation divided by average total assets; return on capital employed (*ROCE*), measured as earnings Before Interest and Taxes divided by average capital employed, where capital employed equals long-term and current debt plus common/ordinary equity; return on equity (*ROE*), computed as net income divided by average book equity, defined as total parent stockholders' equity plus deferred taxes and investment tax credit; and labor expense intensity (*Staff_sale*), measured as labor expenses divided by sales. Our results, reported in Table 10, confirm the consistent impact of AI exposure on organizational structure.

4.2.7 Span of Control

In addition to running various moderation analyses to unpack the mechanisms, we also conduct two additional analyses. Specifically, we first conduct a sub-sample analysis focusing on observations with a span of control greater than zero (i.e., excluding observations with only one seniority level) and one (i.e., excluding observations in which the subordinates have fewer employees, such as the first level). Our results, shown in Table 11, confirm the consistent findings of the AI exposure effect.

In addition, we aggregate the span of control across the lowest four and the highest three seniority levels, aiming to understand how AI exposure has different impacts on frontline versus C-level employees in terms of span of control. Specifically, we designed two dependent variables, *SpanControl1_4* and *SpanControl5_7*, which measure the average span of control across seniority levels 2-4 and 5-7, respectively, representing the span of control of the frontline versus C-level employees. Our results, shown in Table 12, reveal that AI exposure positively increases the span of control for both types of employees, where the impact is bigger among the C-level employees.

4.3 Moderation Mechanisms

We conduct moderation analyses to understand the mechanisms. Specifically, we analyze how employees' technical background, organizational size, the technical industry, and firm age moderate the impact of AI exposure on organizational structure. Building on the DID model framework, we include the moderator into the model and organize the difference-in-difference-in-difference (DDD) model. The specific regression model is shown as equation (2).

$$\{Rank_{it}, SpanControl_{it}\} = \alpha_t + \gamma_i + \zeta_i + \beta_1 AIExposure_{it} * Moderator_{it} + \beta_2 AIExposure_{it} + Control_{it} + \varepsilon_{it}, \quad (2)$$

where we use a time-variant variable, *Moderator_{it}*, to understand the mechanisms of the AI exposure effect. β_1 captures the moderation effects. Other variables are similar to equation (1). Following equation (1), we also run TWFE estimation to understand the moderation effect. To enhance the robustness of our results, we also run an instrumental variable approach for this moderation analysis. We report the results with instrumental variables below but our results without using instrumental variables are consistent and available upon requests.

4.3.1 Employees' Technical Background

We start by examining how employees' technical backgrounds explain the effects of AI exposure. The literature indicates that AI can enhance employees' information-processing capabilities (Jarrahi 2018; Shrestha et al. 2019), thus improving the organizational communication efficiency and empowering managers to supervise more employees. Those can reduce the need for intermediate layers (Bloom et al.

2014). Besides, research suggests individuals with stronger quantitative and technical backgrounds, such as those trained in STEM fields, can be better positioned to leverage AI tools effectively (Bao et al. 2026). Therefore, when a firm has a higher percentage of employees with a technical background, such as a STEM major at a senior level, the effect of AI exposure through improving information processing is more pronounced. Running the moderation analysis, we obtained the results shown in Table 13 (columns 1 and 2), which confirm that employees' technical background enhances the effect of AI exposure, revealing the mechanisms by which AI enhances employees' information processing.

4.3.2 High Tech Industry

We also examine the moderation effect of the tech industry. Consistent with prior research on IT adoption and organizational technology capabilities (Bloom et al. 2014), firms in more technology-intensive sectors are better positioned to adopt and leverage advanced technologies such as AI, thus enhancing the AI exposure impact on the organizational structure. Following the U.S. Census Bureau's NAICS code classification for the high-tech industry (U.S. Census Bureau 2024), we classify each firm as high-tech or non-high-tech. Specifically, we use the following NAICS codes: 3341, 3342, 3344, 3345, 3364, 5112, 5182, 5191, 5413, 5415, and 5417, as listed by the U.S. Census Bureau (2024). We put the description for each NAICS code in Table B.1 of Appendix B. Applying those codes to our data, we found that there are 1,303 organizations in the high-tech industry and 6,378 organizations in the non-high-tech industry. Running the moderation analysis to estimate equation (2) by using the High_Tech_Industry as the moderator, we obtain the results shown in Table 13 (columns 3 and 4), which confirm that the high-tech industry enhances the impact of AI exposure on span of control.

4.3.3 Industry Automation

Industry automation can modify the impact of AI exposure on organizational structure, as AI has been shown to be more effective at handling routine, automated tasks (Wang et al. 2024; Krakowski et al. 2025). To measure industry automation, we follow Afonso and Forte (2024) and assess the level of routine tasks in a particular industry. According to the literature (e.g., Lillrank et al. 2004; Marcolin et al. 2016; Kuntz et al. 2019; Lewandowski et al. 2020), industries like manufacturing, retail, transportation, and

administrative support normally have a high level of routine tasks, while industries like professional services, health care, education, arts, and management normally have a low level of routine tasks. Therefore, we categorize organizations in our data as having high or low routine tasks based on their industry membership. Applying this categorization, we got 3,811 organizations in the high-routine industry and 819 organizations in the low-routine industry. Running the moderation analysis to estimate equation (2) by using the `High_Level_Routine` as the moderator, we obtain the results shown in Table 13 (columns 5 and 6), which reveal that the organization with a high routine task level (vs. a low routine task level) strengthens the impact of AI exposure on span of control.

5 Discussion

Organizational structure plays a central role in shaping how firms coordinate activities, allocate authority, and respond to technological change. As artificial intelligence (AI) diffuses rapidly across industries, executives, consultants, and policymakers increasingly debate whether AI will flatten organizations by reducing managerial layers or, alternatively, introduce new forms of hierarchy and control. Despite the prominence of these debates, systematic evidence on how AI exposure reshapes firms' internal hierarchies has remained limited. This study contributes to the literature by examining the relationship between AI exposure and two fundamental dimensions of organizational structure, hierarchical ranks and span of control—using large-scale longitudinal data from the U.S. labor market.

Leveraging detailed workforce data from Revelio Labs and a staggered difference-in-differences research design, our findings reveal a consistent pattern of organizational restructuring following AI exposure. Specifically, AI exposure is associated with a reduction in the number of hierarchical ranks and an increase in managerial span of control. These results suggest that AI acts as a managerial augmentation technology, enabling firms to compress vertical layers while expanding the scope of oversight for remaining managers. In doing so, AI appears to alter the traditional trade-off between hierarchy depth and managerial capacity by reducing information-processing constraints that historically necessitated additional layers of

management.

Beyond these average effects, our analyses uncover important sources of heterogeneity that shed light on the mechanisms underlying AI-driven organizational change. First, we find that the effects of AI exposure are stronger in firms with a higher concentration of technically trained managers and employees. This pattern suggests that AI's ability to substitute for intermediate managerial functions depends critically on complementary human capital. Managers with stronger technical backgrounds may be better positioned to interpret AI-generated insights, trust algorithmic recommendations, and integrate them into decision-making processes, thereby allowing them to oversee larger teams with fewer hierarchical intermediaries.

Second, organizational size plays a moderating role. Larger firms experience a more pronounced increase in span of control following AI exposure, while the reduction in hierarchical ranks is attenuated. This finding indicates that scale amplifies the coordination benefits of AI but may also impose structural constraints that limit the extent of hierarchy compression. In large organizations, AI appears to enable managerial scaling, allowing firms to expand oversight without proportional increases in headcount—while preserving some vertical differentiation necessary for control and accountability.

Third, we observe stronger effects of AI exposure in technical and highly automated industries. In these settings, AI-related tools are more easily integrated into existing workflows, reinforcing their capacity to standardize decision-making and reduce the need for layered supervision. Together, these findings suggest that AI's organizational consequences are neither uniform nor automatic; rather, they depend on the alignment between technology, human capital, and industry context.

5.1 Literature Contribution and Managerial Implications

This study contributes to multiple streams of literature in information systems, strategy, and organizational design. First, we extend the growing literature on the economic and organizational impacts of AI, which has largely focused on productivity, innovation, and skill demand, by shifting attention to firms' internal structures. While prior research documents how AI changes what tasks are performed and who performs them, our findings demonstrate that AI also reshapes how authority and coordination are organized within firms.

Second, we contribute to the organizational design literature by providing large-scale empirical evidence on how advances in information-processing technologies affect hierarchy and span of control. Building on classic theories that link organizational structure to information-processing constraints, our results show that AI relaxes these constraints in ways that enable flatter hierarchies and broader managerial oversight. In doing so, we provide empirical support for long-standing theoretical claims while highlighting how AI differs from earlier generations of information technology.

Third, our study advances understanding of heterogeneity in technology adoption outcomes. By identifying the moderating roles of technical human capital, firm size, and industry automation, we move beyond average treatment effects to explain when and where AI is most likely to reshape organizational hierarchies. This perspective responds to recent calls for more nuanced analyses of AI adoption that account for organizational and contextual contingencies.

In addition to these academic contributions, our findings offer important managerial implications. For executives and organizational designers, the results suggest that AI adoption decisions cannot be separated from questions of organizational structure. AI exposure enables firms to reduce hierarchical layers and expand spans of control, but these benefits are most pronounced when supported by complementary investments in technical talent and organizational capabilities. Managers should therefore view AI not merely as a productivity-enhancing tool but as a catalyst for structural redesign.

Moreover, our findings caution against one-size-fits-all expectations about AI-driven flattening. In larger firms and certain industries, hierarchy reduction may be limited even as managerial spans expand. Leaders should anticipate these differential effects and tailor restructuring efforts accordingly. For example, firms may need to invest in training managers to work effectively with AI systems, redesign reporting relationships, or recalibrate performance evaluation processes to sustain effectiveness under wider spans of control.

Overall, our study provides evidence-based guidance to inform ongoing industry debates about the future of managerial work and organizational hierarchy in the AI era. By clarifying how AI exposure interacts with organizational context to shape structure, we help managers move beyond speculative claims

toward more informed and strategic decision-making.

5.2 Limitations and Future Research

As with any empirical study, our findings should be interpreted in light of several limitations that point to opportunities for future research. First, our analysis focuses on U.S. firms, which may limit the generalizability of the results to other institutional contexts. Future studies could examine whether similar patterns emerge in countries with different labor market institutions or managerial norms.

Second, while our AI exposure measure captures the presence of AI-related human capital within firms, it does not directly observe how AI tools are deployed in practice. Future research could combine workforce data with information on internal AI applications or investments to better distinguish between different modes of AI use.

Third, our measures of hierarchy and span of control are inferred from job titles and predicted seniority levels. Although this approach follows established practice in the literature, future work could validate these measures using internal organizational charts or survey data where available.

Fourth, our analysis focuses on medium-term structural changes and may not fully capture longer-term dynamics. Organizational restructuring in response to AI may unfold gradually as firms experiment, learn, and adjust. Longitudinal studies with extended post-adoption windows could shed light on the persistence or evolution of AI-driven structural changes.

Finally, future research could explore downstream consequences of AI-induced organizational restructuring, such as impacts on employee performance, career mobility, wage inequality, or firm resilience. Understanding these outcomes would further illuminate the broader implications of AI for organizations and labor markets.

References

- Acemoglu D, Restrepo P (2018) Artificial intelligence, automation, and work. In *The economics of Artificial Intelligence: An Agenda* (pp. 197-236). University of Chicago Press.
- Afonso Ó, Forte R (2024) Routine and non-routine sectors, tasks automation and wage polarization. *Applied Economics*. 56(55):7262-7285.

- Agarwal S, Qian W, Zou X (2021) Disaggregated sales and stock returns. *Management Science*. 67(11):7167-7183.
- Alyakoob M, Rahman MS (2022) Shared prosperity (or lack thereof) in the sharing economy. *Information Systems Research*. 33(2):638-658.
- Angrist JD, Pischke JS (2009) *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Anthony C, Bechky BA, Fayard AL (2023) “Collaborating” with AI: Taking a system view to explore the future of work. *Organization Science*. 34(5):1672-1694.
- Babina T, He AX, Howell ST, Perlman ER, Staudt J (2023) Cutting the innovation engine: how federal funding shocks affect university patenting, entrepreneurship, and publications. *The Quarterly Journal of Economics*. 138(2), 895-954.
- Babina T, Fedyk A, He A, Hodson J (2024) Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*. 151:103745.
- Baker AC, Larcker DF, Wang CC (2022) How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*. 144(2):370-395.
- Bao L, Huang D, Lin C (2026) Can artificial intelligence improve gender equality? Evidence from a natural experiment. *Management Science*. 72(1):474-494.
- Berente N, Gu B, Recker J, Santhanam R (2021) Managing artificial intelligence. *MIS Quarterly*. 45(3):1433-1450.
- Bloom N, Garicano L, Sadun R, Van Reenen J (2014) The distinct effects of information technology and communication technology on firm organization. *Management Science*. 60(12):2859-2885.
- Brynjolfsson E, Li D, Raymond L (2025) Generative AI at work. *The Quarterly Journal of Economics*. 140(2):889-942.
- Burtch G, Carnahan S, Greenwood BN (2018) Can you gig it? An empirical examination of the gig economy and entrepreneurial activity. *Management Science*. 64(12):5497-5520.
- Cengiz D, Dube A, Lindner A, Zipperer B (2019) The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*. 134(3):1405-1454.
- Cheng Z, Rai A, Tian F, Xu SX (2021) Social learning in information technology investment: The role of board interlocks. *Management Science*. 67(1):547-576.
- Cheng HK, Zhao X, Li Y, Wan XS, Zang X (2025) Bridge of cryptos: The impact of cross-chain technology on market efficiency. SSRN 5285744.
- Child J (1987) Information technology, organization, and the response to strategic challenges. *California Management Review*. 30(1):33-50.
- Cohen LE, Mahabadi S (2022) In the midst of hiring: Pathways of anticipated and accidental job evolution during hiring. *Organization Science*. 33(5):1938-1963.
- Colombo M, Delmastro M (2008) *The economics of organizational design: Theoretical insights and empirical evidence*. Springer.
- Davenport T, Guha A, Grewal D, Bressgott T (2020) How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*. 48(1):24-42.
- De Chaisemartin, C., & d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, 1-45.
- Fountain T, McCarthy B, Saleh T (2019) Building the AI-powered organization. *Harvard Business Review*. 97(4):62-73.
- Galbraith JR (1974) Organization design: An information processing view. *Interfaces*. 4(3):28-36.
- Gavett G, Sawhney V (2025). What’s the future of middle management? <https://hbr.org/2025/04/whats-the-future-of-middle-management>.
- Gómez J, Salazar I, Vargas P (2016) Firm boundaries, information processing capacity, and performance in manufacturing firms. *Journal of Management Information Systems*. 33(3):809-842.
- Goodman-Bacon A (2021) Difference-in-differences with variation in treatment timing. *Journal of*

- Econometrics. 225(2): 254-277.
- Guadalupe M, Li H, Wulf J (2014) Who lives in the C-suite? Organizational structure and the division of labor in top management. *Management Science*. 60(4):824-844.
- Huang S, Li S, Wang R (2025) From structure to quality: How hierarchies in accounting function shape financial reporting outcomes. Available at SSRN 5282420.
- Imbens, GW, Angrist JD (1994) Identification and estimation of local average treatment effects. *Econometrica*. 62(2):467-475.
- Jarrahi MH (2018) Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*. 61(4):577-586.
- Keren M, Levhari D (1979) The optimum span of control in a pure hierarchy. *Management Science*. 25(11):1162-1172.
- Krakowski S, Haftor D, Luger J, Pashkevich N, Raisch S (2025) Human-centered artificial intelligence: A field experiment. *Management Science*.
- Kuntz L, Scholtes S, Sülz S (2019) Separate and concentrate: Accounting for patient complexity in general hospitals. *Management Science*. 65(6):2482-2501.
- Lewandowski P, Park A, Schotte S (2020) The global distribution of routine and non-routine work (No. 13384). Institute of Labor Economics (IZA).
- Li Q, Lourie B, Nekrasov A, Shevlin T (2022) Employee turnover and firm performance: Large-sample archival evidence. *Management Science*. 68(8):5667–5683.
- Liang C, Lourie B, Nekrasov A (2024) Does the risk aversion of accountants matter? Female rank-and-file accounting employees and internal control quality. *Journal of Accounting and Public Policy*. 45:107194.
- Liang C, Lourie B, Nekrasov A, Shevlin T (2025) The gender position gap and firm performance. *Journal of Business Finance & Accounting*. 52(5):2464-2491.
- Lillrank P, Liukko M (2004) Standard, routine and non-routine processes in health care. *International Journal of Health Care Quality Assurance*. 17(1):39-46.
- Liu L, Wang Y, Xu Y (2024) A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. *American Journal of Political Science*. 68(1):160-176.
- Lou B, Ma X, Wu L (2025) Artificial intelligence, CEO turnover, and exploration orientation in firm innovation. *Information Systems Research*.
- Lou B, Wu L (2021) AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms. *MIS Quarterly*. 45(3):1451-1482.
- Lu LX, Lu SF (2022) Does nonprofit ownership matter for firm performance? Financial distress and ownership conversion of nursing homes. *Management Science*. 68(7):5127-5145.
- Marchetti A, Puranam P (2025) Are less hierarchical firms organized around stronger cultures? Evidence from big data. *Strategic Management Journal*. Forthcoming.
- Marcolin L, Miroudot S, Squicciarini M (2016) Routine jobs, employment and technological innovation in global value chains (No. 2016/1). Paris: OECD Publishing.
- Marsh RM, Mannari H (1981) Technology and size as determinants of the organizational structure of Japanese factories. *Administrative Science Quarterly*. 33-57.
- McElheran K, Li JF, Brynjolfsson E, Kroff Z, Dinlersoz E, Foster L, Zolas N (2024) AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*. 33(2):375-415.
- McKinsey Global Institute (2023) The economic potential of generative AI: The next productivity frontier. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>.
- Mihm J, Loch CH, Wilkinson D, Huberman BA (2010) Hierarchical structure and search in complex organizations. *Management Science*. 56(5):831-848.
- Nolan B (2025) AI is already upending the corporate org chart as it flattens the distance between the C-suite and everyone else. <https://finance.yahoo.com/news/ai-already-upending-corporate-org-083100729.html>.

- Pu J, Kwark Y, Han SP, Ye Q, Gu B (2024) Uncertainty reduction vs. reciprocity: Understanding the effect of a platform-initiated reviewer incentive program on regular ratings. *Information Systems Research*. 35(3):1363-1381.
- Qiao D, Rui H, Xiong Q (2026) Artificial intelligence and jobs: Has the inflection point arrived? Evidence from an online labor platform. *Journal of Management Information Systems*. 43(1):5-37.
- Rai A, Tian J, Xue L (2026) FAIR: A design theory for artificial intelligence fairness. *MIS Quarterly*.
- Raisch S, Krakowski S (2021) Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*. 46(1):192-210.
- Rajan RG, Wulf J (2006) The flattening firm: Evidence from panel data on the changing nature of corporate hierarchies. *The Review of Economics and Statistics*. 88(4):759-773.
- Revelio Labs (2024) <https://www.data-dictionary.reveliolabs.com/methodology.html#jobs-taxonomy>.
- Rothaermel FT (2018) *Strategic management*. New York, NY: McGraw-Hill.
- Shrestha YR, Ben-Menahem SM, Von Krogh G (2019) Organizational decision-making structures in the age of artificial intelligence. *California Management Review*. 61(4):66-83.
- Simon HA (1960) *The new science of management decision*. Harper & Brothers.
- Spring M, Faulconbridge J, Sarwar A (2022) How information technology automates and augments processes: Insights from Artificial-Intelligence-based systems in professional service operations. *Journal of Operations Management*. 68(6-7):592-618.
- Stock JH, Yogo M (2005) Testing for weak instruments in Linear IV regression. In D. W. K. Andrews & J. H. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg* (pp. 80-108). Cambridge University Press.
- Sull D, Sull C, Chamberlin A (2019) Measuring culture in leading companies. *MIT Sloan Management Review*.
- Sull D, Sull C, Zweig B (2022) Toxic culture is driving the great resignation. *MIT Sloan Management Review*.
- Tambe PB (2026) Reskilling the workforce for AI: Domain expertise and algorithmic literacy. *Management Science*. 72(1):515-537.
- Tschang FT, Almirall E (2021) Artificial intelligence as augmenting automation: Implications for employment. *Academy of Management Perspectives*. 35(4):642-659.
- Tremblay D (2025) How AI is reshaping organizational design: Global implementation trends. <https://www.hireborderless.com/post/how-ai-is-reshaping-organizational-design>.
- Trustible (2024) Why AI Governance is going to get a lot harder. <https://trustible.ai/post/why-ai-governance-is-going-to-get-a-lot-harder/>.
- U.S. Census Bureau (2024) BDS-High tech methodology. <https://www.census.gov/programs-surveys/ces/data/public-use-data/experimental-bds/bds-high-tech/methodology.html>.
- Vergne JP (2020) Decentralized vs. distributed organization: Blockchain, machine learning and the future of the digital platform. *Organization Theory*, 1(4): 2631787720977052.
- Wang J, Li G, Hui KL (2022) Monetary incentives and knowledge spillover: Evidence from a natural experiment. *Management Science*. 68(5):3549-3572.
- Wang L, Huang N, Hong Y, Liu L, Guo X, Chen G (2023) Voice-based AI in call center customer service: A natural field experiment. *Production and Operations Management*. 32(4):1002-1018.
- Wang W, Gao G, Agarwal R (2024) Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science*. 70(9):5753-5775.
- Wang L, Huang N, He Y, Liu D, Guo X, Sun Y, Chen G (2025) Artificial intelligence (AI) assistant in online shopping: A randomized field experiment on a livestream selling platform. *Information Systems Research*. 36(4):2358-2374.
- Wang N, Li Z (2025) Channel structures for a dual-purpose firm: Direct, retail, or hybrid? *Production and Operations Management*. 34(8):2288-2302.
- Wang X, Wu L (2026) Artificial intelligence, lean startup method, and product innovations. *Management*

Science. 72(1):756-782.

Wei Z, Lin M (2017) Market mechanisms in online peer-to-peer lending. Management Science. 63(12):4236-4257.

Wu L, Lou B, Hitt LM (2025) Innovation strategy after IPO: How AI analytics spurs innovation after IPO. Management Science. 71(3):2360-2389.

Xue M, Cao X, Feng X, Gu B, Zhang Y (2022) Is college education less necessary with AI? Evidence from firm-level labor structure changes. Journal of Management Information Systems. 39(3):865-905.




Figure 1. Parallel Trend

Table 1. Summary Statistics

Variables	Observations	Mean	Std.Dev.	Min	Max
AIExposure	171836	.095	.293	0	1
Rank	171836	6.333	1.303	1	7
SpanControl	171836	3.33	2.193	0	104.647
Log(Size)	171836	3.649	2.268	.693	13.19

Note. For span of control, around 2% of observations have a span of control of less than 1, where the first seniority level is internships, which are normally smaller than the second seniority level. In addition, some firms have only one rank in some years, particularly in their starting years, resulting in a span of control of 0 and accounting for around 1% of observations. We conducted robustness checks by removing those observations and obtaining consistent results. More details are presented in the robustness check section.

Table 2. Estimation Results of the DID Model using TWFE

	(1) Rank	(2) SpanControl
AIExposure	-0.203*** (0.01)	0.0278* (0.02)
Log(Size)	0.192*** (0.00)	0.0471*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1122	0.0037

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format.

Table 3. IV Approach Main Effects

Variable	(0) AIExposure	(1) Rank	(2) SpanControl
	1 st stage of 2SLS	2 nd stage of 2SLS	
AITalents	0.006*** (0.001)		
AIExposure		-0.562*** (0.06)	0.573*** (0.13)
Log(Size)	0.054*** (0.0003)	0.215*** (0.00)	0.0164** (0.01)
Firm Fixed Effect	/	Yes	Yes
Year Fixed Effect	/	Yes	Yes
Observations	171,836	171805	171805
R ²	0.19 (F-stat = 10504.62)	0.1081	0.0038

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format. The independent variable, AIExposure is instrumented using AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities.

Table 4. Falsification Test

	(1) Rank	(2) SpanControl	(3) Rank	(4) SpanControl
μ of Random β	0.000219	-0.000057	0.000221	-0.000056
σ of Random β	0.000127	0.000298	0.000127	0.000298
Replications	1000	1000	1000	1000
Estimated β	-0.203	0.0278	-0.203	0.0278
Z-Score	1602.109	-93.445	1601.653	-93.414
P-value	.000	.000	.000	.000

Note. The first two columns used the organization and year fixed effect, while the last two columns used the firm, year, and state fixed effects. μ and σ are the mean and standard error of the placebo estimators. This diagnostic test is to determine the probability of the observed effect occurring purely by chance. The comparison between a placebo effect and an estimated effect is statistically significantly different ($p < 0.01$), thus eliminating the aforementioned possibility.

Table 5. Estimation Results of IFect Estimator

	(1) Rank	(2) SpanControl
AIExposure	-0.119*** (0.02)	0.081* (0.04)
Log(Size)	0.212*** (0.01)	0.051*** (0.02)
Observation	171836	171836

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format.

Table 6. Estimation Results Using Alternative AI Exposure Measurement

	(1) Rank	(2) SpanControl
AIExposure_2	-0.147*** (0.01)	0.105*** (0.01)
Log(Size)	0.192*** (0.00)	0.0425*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R^2	0.1121	0.0042

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. AIExposure_2 is measured by the number of AI workers a firm hires in a given year, with a log transformation. Size measures the number of employees an organization has in a given year, in log-transformed form.

Table 7. Estimation Results using Alternative Outcome Variables and Models

Panel A: Log-transformed Outcome Variables		
	(1)	(2)
	Log(Rank)	Log(SpanControl)
AIExposure	-0.0242*** (0.00)	0.105*** (0.01)
Log(Size)	0.0312*** (0.00)	0.0425*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.0810	0.0042

Panel B: Poisson Model		
	(1)	(2)
	Rank	SpanControl
AIExposure	-0.0230*** (0.00)	0.0288*** (0.01)
Log(Size)	0.0298*** (0.00)	0.0135*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171760	170769
Wald Chi2	811.75	258.79

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format. In the fixed-effects Poisson models, Stata drops panels with insufficient within-unit variation by construction, thereby reducing the final sample size.

Table 8. Estimation Results Using Lagging Dependent Variables

	Lag One-year		Lag Two-year		Lag Three-year	
	(1) Rank	(2) SpanControl	(1) Rank	(2) SpanControl	(1) Rank	(2) SpanControl
AIExposure	-0.0277*** (0.01)	0.00895*** (0.00)	-0.0211*** (0.01)	0.00844** (0.00)	-0.0176** (0.01)	0.00850** (0.00)
Log(Size)	0.0656*** (0.00)	0.00854*** (0.00)	0.0621*** (0.00)	0.00529** (0.00)	0.0594*** (0.00)	0.00317 (0.00)
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44961	44961	41211	41211	37640	37640
R ²	0.0173	0.0012	0.0161	0.0007	0.0146	0.0004

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format.

Table 9. Estimation Results using Lagging and Removing Company Size

Panel A: Lagging the Company Size		
	(1)	(2)
	Log(Rank)	Log(SpanControl)
AIExposure	-0.188*** (0.01)	0.0311** (0.02)
Size_lag1	0.178*** (0.00)	0.0480*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	163450	163450
R ²	0.0974	0.0039

Panel B: Removing the Control of Company Size		
	(1)	(2)
	Rank	SpanControl
AIExposure	-0.0656*** (0.01)	0.0615*** (0.02)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.0378	0.0027

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size_lag1 measures the lagged one year of the number of employees an organization has in a particular year in its log transformation format. Poisson models, Stata drops panels with insufficient within-unit variation by construction, thereby reducing the final sample size.

Table 10. Estimation Results Including Financial Performance Controls

	(1)	(2)
	Rank	SpanControl
AIExposure	-0.105*** (0.01)	0.219*** (0.03)
Log(Size)	0.0973*** (0.00)	0.169*** (0.01)
Adv_sale	-0.0162 (0.01)	-0.0573 (0.04)
Capital_ratio	0.000254 (0.01)	-0.0154 (0.02)
Cash_ratio	0.00118 (0.00)	-0.00125 (0.00)
ROA	-0.0130 (0.01)	0.0603* (0.04)
ROCE	0.000553*** (0.00)	-0.0000370 (0.00)
ROE	0.0000176 (0.00)	0.0000673 (0.00)
Staff_sale	0.0000676 (0.00)	0.000228 (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	29755	29755
R ²	0.0884	0.0212

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format. The number of observations is reduced because CompuStat only provides financial performance for public firms.

Table 11. Estimation Results Among the Sub-samples

Panel A. Span of Control Larger than Zero		
	(1)	(2)
	Rank	SpanControl
AIExposure	-0.196*** (0.01)	0.0355** (0.02)
Log(Size)	0.185*** (0.00)	0.0434*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	168760	168760
R ²	0.1155	0.0039
Panel B. Span of Control Larger than One		
	(1)	(2)
	Rank	SpanControl
AIExposure	-0.156*** (0.01)	0.0660*** (0.02)
Log(Size)	0.179*** (0.00)	0.0624*** (0.00)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	151948	151948
R ²	0.1032	0.0044

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format.

Table 12. Estimation Results on Span of Control for Frontier versus C-level Employees

	(1)	(2)
	SpanControl1 4	SpanControl5 7
AIExposure	0.0137** (0.01)	0.333*** (0.03)
Log(Size)	-0.0733*** (0.00)	-0.000859 (0.01)
Firm Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	153465	120016
R ²	0.0267	0.0078

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Size measures the number of employees an organization has in a particular year in its log transformation format.

Table 13. Moderation Mechanism of Employees' Technical Background (Instrumented)

	(1) Rank	(2) SpanControl	(3) Rank	(4) SpanControl	(5) Rank	(6) SpanControl
AIExposure *	-5.419*** (0.60)	5.704*** (1.34)				
S7_stem_ratio						
AIExposure *			-0.969*** (0.11)	1.049*** (0.25)		
High_Tech_Industry						
AIExposure *					-0.497*** (0.06)	0.538*** (0.13)
High_Level_Routine						
AIExposure	-0.205*** (0.01)	0.0268* (0.02)	-0.202*** (0.01)	0.0264* (0.02)	-0.202*** (0.01)	0.0264* (0.02)
S7_stem_ratio	0.512*** (0.02)	-0.0657 (0.05)				
Log(Size)	0.190*** (0.00)	0.0469*** (0.00)	0.192*** (0.00)	0.0466*** (0.00)	0.192*** (0.00)	0.0466*** (0.00)
Firm Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes
Observations	171836	171836	171836	171836	171836	171836
R ²	0.1157	0.0038	0.1126	0.0038	0.1126	0.0038

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. We used an instrumental variable approach, where both the AIExposure and the interaction of it with a moderator are instrumented. Specifically, AIExposure is instrumented using AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator. S7_stem_ratio measures the ratio of employees with a STEM education at the seniority level 7. We also tried other ratios of employees with a STEM education at other seniority levels, but got consistent results, as shown in Tables A.1-A.6 of Appendix A. Size measures the number of employees an organization has in a particular year in its log transformation format. High_Tech_Industry measures whether a firm is in the high-tech industry. The coefficients associated with it were omitted because of the organization fixed effect. High-Level_Routine measures whether the organization belongs to the industry with a high-level routine or a low-level routine. Specifically, high_level_routine=1 is an industry of manufacturing, retail, transportation, or administrative support, and high_level_routine=0 is an industry of professional services, health care, education, arts, or management.

Appendix

Appendix A. Additional Results

Table A.1. Moderation Mechanism of Employees' Technical Background (Seniority = 6)

	(1) Rank	(2) SpanControl
AIExposure * S6_stem_ratio	-2.490*** (0.26)	2.430*** (0.57)
AIExposure	-0.213*** (0.01)	0.0266* (0.02)
S6_stem_ratio	0.649*** (0.02)	-0.00934 (0.04)
Log(Size)	0.186*** (0.00)	0.0467*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1202	0.0038

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S6_stem_ratio measures the ratio of employees with a STEM education at the seniority level 6. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S6_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Table A.2. Moderation Mechanism of Employees' Technical Background (Seniority = 5)

	(1) Rank	(2) SpanControl
AIExposure * S5_stem_ratio	-2.416*** (0.23)	1.972*** (0.52)
AIExposure	-0.225*** (0.01)	0.0138 (0.02)
S5_stem_ratio	1.064*** (0.03)	0.581*** (0.06)
Log(Size)	0.181*** (0.00)	0.0406*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1219	0.0044

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S5_stem_ratio measures the ratio of employees with a STEM education at the seniority level 5. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S5_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Table A.3. Moderation Mechanism of Employees' Technical Background (Seniority = 4)

	(1) Rank	(2) SpanControl
AIExposure * S4_stem_ratio	-2.148*** (0.21)	1.852*** (0.47)
AIExposure	-0.223*** (0.01)	0.0157 (0.02)
S4_stem_ratio	0.845*** (0.02)	0.419*** (0.05)
Log(Size)	0.181*** (0.00)	0.0410*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1205	0.0042

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S4_stem_ratio measures the ratio of employees with a STEM education at the seniority level 4. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S4_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Table A.4. Moderation Mechanism of Employees' Technical Background (Seniority = 3)

	(1) Rank	(2) SpanControl
AIExposure * S3_stem_ratio	-2.061*** (0.20)	1.819*** (0.45)
AIExposure	-0.227*** (0.01)	0.0171 (0.02)
S3_stem_ratio	0.843*** (0.02)	0.308*** (0.04)
Log(Size)	0.179*** (0.00)	0.0418*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1239	0.0041

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S3_stem_ratio measures the ratio of employees with a STEM education at the seniority level 3. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S3_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Table A.5. Moderation Mechanism of Employees' Technical Background (Seniority = 2)

	(1) Rank	(2) SpanControl
AIExposure * S2_stem_ratio	-1.836*** (0.16)	1.272*** (0.37)
AIExposure	-0.244*** (0.01)	-0.00183 (0.02)
S2_stem_ratio	1.226*** (0.02)	0.825*** (0.04)
Log(Size)	0.167*** (0.00)	0.0298*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1339	0.0060

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S2_stem_ratio measures the ratio of employees with a STEM education at the seniority level 2. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S2_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Table A.6. Moderation Mechanism of Employees' Technical Background (Seniority = 1)

	(1) Rank	(2) SpanControl
AIExposure * S1_stem_ratio	-2.802*** (0.28)	2.396*** (0.63)
AIExposure	-0.220*** (0.01)	0.0137 (0.02)
S1_stem_ratio	1.203*** (0.02)	0.844*** (0.05)
Log(Size)	0.177*** (0.00)	0.0356*** (0.00)
Organization Fixed Effect	Yes	Yes
Year Fixed Effect	Yes	Yes
Observations	171836	171836
R ²	0.1260	0.0053

Note. Robust standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. S1_stem_ratio measures the ratio of employees with a STEM education at the seniority level 1. Size measures the number of employees an organization has in a particular year in its log transformation format. We used an instrumental variable approach, where both the AIExposure and AIExposure * S1_stem_ratio are instrumented. Specifically, AIExposure is instrumented using the AITalents, which measures firms' exposure to the supply of AI talent from U.S. universities. The interaction is instrumented using the interaction between AITalents and the moderator.

Appendix B. NACIS Codes of High-Tech Industry

Table B.1. Description of High-Tech Industry NACIS Codes

NAICS	Description
3341	Computer and Peripheral Equipment Manufacturing
3342	Communications Equipment Manufacturing
3344	Semiconductor and Other Electronic Component Manufacturing
3345	Navigational, Measuring, Electromedical, and Control Instruments Manufacturing
3364	Aerospace Product and Parts Manufacturing
5112	Software Publishers
5182	Data Processing, Hosting, and Related Services
5191	Other Information Services
5413	Architectural, Engineering, and Related Services
5415	Computer Systems Design and Related Services
5417	Scientific Research and Development Services