How AI Shocks Dissipate Executive Rents and Reshape Compensation Stickiness *

Jianwei Chen† **Jialiang He**‡ **Xinyan Liu**§ **Dan Wang**ξ

Abstract

Compensation stickiness has two classical rationales: implicit insurance for risk-averse employees (Baily, 1974; Azariadis, 1975) and managerial power for executives (Bebchuk and Fried, 2003). However, Artificial Intelligence (AI), by reshaping task allocation and information flows within firms, disrupts both. We construct an analytical model integrating knowledge hierarchy and task automation frameworks. Using China's National New Generation AI Innovation and Development Pilot Zones as a quasi-natural experiment, we apply a heterogeneity-robust imputation estimator to a panel of A-share listed firms from 2015 to 2024 for causal identification. We find that AI shocks induce asymmetric compensation stickiness within firms: executive compensation exhibits substantial downward flexibility (by approximately 6.1%), whereas the wages of rank-and-file employees maintain stable downward rigidity, narrowing the intra-firm pay gap. Mechanism analyses indicate that AI adoption drives executive rent dissipation—manifested as dual declines in excess compensation and excess perquisites—which serves as the primary transmission channel for this evolving compensation landscape. Heterogeneity tests reveal that this effect is more pronounced in firms with high internal information asymmetry, weak managerial outside options, high labor intensity, and those in their growth stages. This study extends the intersection of technological change and sticky wage theories from a vertical hierarchy perspective.

Keywords: AI; Rent Dissipation; Compensation Stickiness; Intra-firm Pay Gap

JEL Classification: J31; O33

1 Introduction

* We would like to thank Professor Lifeng Su (University of International Business and Economics) and Associate Professor Xiaobo He (Shanghai University of International Business and Economics) for their valuable comments and suggestions on this paper.

† University of International Business and Economics. Email: chenjianwei@uibe.edu.cn

‡ University of International Business and Economics. Email: hjl6424@163.com

§ University of International Business and Economics. Email: liuxy_689@163.com Corresponding author.

ξ University of International Business and Economics. Email: 18702600458@163.com

Major technological shocks reshape both aggregate growth and the distribution of its gains. During Britain's First Industrial Revolution, for example, per capita output rose while real wages stagnated — the "Engels' pause" (Allen, 2009). Today, AI raises similar distributional concerns. Existing studies examine AI through a macroeconomic lens, focusing on factor income shares (Guo, 2019; Acemoglu and Restrepo, 2026), skill premiums (Jiang et al., 2025; Bloom et al., 2025), aggregate wealth inequality (Yang and Zhen, 2026), and wage markdowns (Zhu et al., 2026). These studies capture the macro-distributional effects of AI, but they overlook intra-firm power dynamics, task reallocation, and rent sharing. How does AI disrupt existing rent-sharing arrangements? And how does it alter the intra-firm pay gap across corporate hierarchies? These questions remain open.

Compensation stickiness reveals how firms and workers share rents through implicit contracts. In the classical view, downward wage rigidity is an implicit insurance contract: firms shield risk-averse workers from wage cuts in exchange for lower average pay (Baily, 1974; Azariadis, 1975). Within this framework, firms implicitly commit to smoothing compensation during economic fluctuations, thereby endogenizing wage stability as a fundamental risk-sharing arrangement. Extending this analysis to hierarchical disparities, prior studies highlight that executives, leveraging managerial power and substantive control over organizational resources, are better positioned than rank-and-file employees to maintain downward rigidity in their own pay during corporate downturns (Bebchuk and Fried, 2003). Consequently, an asymmetric dichotomy emerges: rank-and-file employees are shielded by implicit insurance contracts, whereas executives are insulated by managerial power. Together, these dual mechanisms constitute the traditional institutional foundation for shielding compensation structures from adverse shocks—a baseline that AI shocks are poised to profoundly disrupt.

However, technological change, typified by AI, fundamentally challenges this paradigm. First, as a general-purpose technology characterized by capacities for rapid evolution and continuous improvement (Baily et al., 2026), intelligent capital exhibits rapid endogenous depreciation. Continuous advancements in AI necessitate large-scale

capital investments, accelerating the obsolescence of existing capital stocks (Zhang and Zhang, 2026). This dynamic diverts corporate cash flows toward rapidly depreciating computational assets, thereby compressing the pool of distributable surplus. Second, AI is reshaping task allocation boundaries within the firm's production process, theoretically leading to the asymmetric dissipation of task-specific rents (Acemoglu and Restrepo, 2026). Consequently, AI-driven capital accumulation and task restructuring may erode the bargaining power of specific employee cohorts, ultimately culminating in positional rent dissipation.

As AI shocks drive the reallocation of intra-firm tasks, a critical question emerges: does the compensation of executives and rank-and-file employees maintain its historical stickiness? Furthermore, do these shocks induce asymmetric compensation stickiness across hierarchical levels? The extant literature has yet to adequately address these questions. In the age of AI, understanding the evolving landscape of compensation stickiness carries profound macroeconomic implications. As global policymakers deploy active fiscal and monetary interventions to channel resources toward digital and intelligent transitions, the distributional mechanisms of the labor market are fundamentally reshaped. If AI shocks dismantle traditional implicit contracts within micro-level firms, established inflation transmission channels and employment stabilization policies may be rendered ineffective. Therefore, elucidating how AI shocks reshape intra-firm compensation dynamics is not merely a matter of corporate governance; it is essential for comprehending the broader evolution of income distribution under the paradigm of advanced productivity.

This paper posits that AI shocks drive the restructuring of task allocation boundaries and information environments across corporate hierarchies, triggering a structural dissipation of task-specific rents and leading to asymmetric compensation stickiness. To elucidate the underlying mechanism, we construct an analytical model integrating knowledge hierarchy and task automation frameworks. Two primary mechanisms operate within this model. First, AI reduces the cost of acquiring explicit knowledge, thereby pushing tasks originally reserved for executives down to the frontlines of production. As rank-and-file employees benefit from reduced learning

costs and assume expanded roles with increased task complexity, their wages exhibit stable downward rigidity. Second, this downward delegation of tasks endogenously compresses the executive span of control, eliminating the bureaucratic rents attached to complex managerial chains. We therefore hypothesize that, under AI shocks, executive compensation is stripped of its traditional institutional insulation and exhibits pronounced downward flexibility. This asymmetry—intertwining downward flexibility in executive compensation with downward stickiness for rank-and-file workers—constitutes a novel feature of intra-firm income distribution in the AI era.

To empirically test these hypotheses, we exploit China's National New Generation AI Innovation and Development Pilot Zones as a quasi-natural experiment. For causal identification, we apply a heterogeneity-robust imputation estimator to a panel of A-share listed firms from 2015 to 2024. Our baseline estimates indicate that AI shocks significantly narrow the intra-firm pay gap. These baseline results remain robust after accounting for multiple potential confounders, including macroeconomic shocks, concurrent policy overlaps, micro-level demand shifts, and major public health emergencies. Furthermore, disaggregated analyses confirm the presence of asymmetric compensation stickiness: AI shocks induce a significant decrease in executive compensation (by approximately 6.1%), whereas the wages of rank-and-file employees maintain stable downward rigidity, empirically validating our theoretical predictions. Mechanism analyses indicate that AI adoption exacerbates executive rent dissipation—manifested as dual declines in excess compensation and excess perquisites—which serves as the primary transmission channel for this evolving compensation landscape.

We contribute to the extant literature in three primary ways. First, we advance the theoretical understanding of technological shocks by incorporating organizational hierarchy into the task rent dissipation framework, illustrating how these shocks yield asymmetric effects across managerial and non-managerial echelons. Previous research predominantly focuses on horizontal task substitution, exploring how technological shocks encroach upon workers' task spaces to trigger job rent dissipation (Wang et al., 2022; Chen et al., 2024; Yu et al., 2025). Less attention has been devoted to heterogeneous responses along the vertical hierarchy of corporate organizations. For

instance, while Acemoglu and Restrepo (2026) demonstrate that automation technologies directly substitute for high-rent tasks, their analytical framework relies on a predominantly parallel task space, primarily addressing changes in aggregate task shares between automation and labor within the macro labor market. We extend the logic of rent dissipation to the intra-firm vertical hierarchy, demonstrating that AI technology can deconstruct the positional rents executives derive from their span of control.

Second, we provide novel empirical evidence on hierarchical wage differentials among Chinese A-share listed firms, thereby enriching the literature on AI technology and the intra-firm pay gap. Existing literature typically emphasizes relative productivity shifts and skill premiums within the macroeconomic labor market (Chen and Qin, 2022; Jiang et al., 2025; Brynjolfsson et al., 2025; Yang and Zhen, 2026), or analyzes the impact of AI on income through the lens of AI-driven factor income shares (Guo, 2019; He et al., 2023; Su et al., 2025; Yu et al., 2025; Zhu et al., 2026). Exploiting the quasi-natural experiment of the National New Generation AI Innovation and Development Pilot Zones, our study offers robust causal identification from the vantage point of corporate organizational hierarchies. We not only document the narrowing effect of AI shocks on the intra-firm pay gap but also uncover the asymmetric compensation stickiness inherent in wage dynamics through a granular examination of executive versus rank-and-file pay. Furthermore, by introducing proxies such as excess compensation and excess perquisites, alongside heterogeneity tests on internal information asymmetry, managerial outside options, labor intensity, and growth stages, we provide robust micro-data support for the rent dissipation mechanism, effectively augmenting the literature on AI-driven organizational restructuring.

Finally, our findings offer critical insights into mitigating the risk of an "Engels' pause" in the age of AI. We reveal that the narrowing of the intra-firm pay gap triggered by AI shocks is driven not by rapid, productivity-enhancing wage increases for rank-and-file employees, but by the partial dissipation of managerial task rents. This suggests that the deployment of AI technology transcends mere efficiency gains; it can also serve as a governance mechanism to deconstruct bureaucratic rent-seeking environments and

optimize intra-firm distribution structures. Concurrently, we highlight that AI imparts an exceptionally high endogenous depreciation rate to digital and intelligent capital. The consequent acceleration of massive capital expenditures may structurally crowd out the capacity for wage growth, potentially precipitating a distributional dilemma wherein output growth coexists with wage stagnation, offering insights into fostering advanced productivity while balancing efficiency and equity.

We develop the model in Section 2, describe the data and identification strategy in Section 3, present baseline results in Section 4, explore mechanisms in Section 5, and conclude in Section 6.

2 Analytical Framework and Theoretical Hypotheses

We model how AI reshapes tasks, rents, and compensation within a knowledge hierarchy. The framework combines Garicano's (2000) theory of knowledge-based hierarchies with Acemoglu and Restrepo's (2018) task automation framework. Building upon the knowledge hierarchy theory established by Garicano (2000), we develop a baseline model that abstracts the production process as a problem-solving endeavor wherein different hierarchical echelons collaborate through knowledge application. We subsequently incorporate the task automation framework of Acemoglu and Restrepo (2018). By modeling the problems to be solved as a continuous distribution based on complexity, we characterize the reshaping effect of AI technology—conceptualized as capital embodying digitalization and intelligence—on intra-firm task boundaries. Distinct from traditional technological progress, we emphasize that AI possesses the unique attribute of reducing the costs associated with acquiring explicit knowledge. The model's derivations demonstrate how AI drives shifts in task boundaries across corporate hierarchies and precipitates the dissipation of executive rents, thereby providing theoretical micro-foundations for our subsequent empirical tests on asymmetric compensation stickiness.

2.1 Representative Firm and Task Allocation

Consider a representative firm organized as a knowledge hierarchy. A unit measure of tasks spans a continuum of complexity $\theta \in [0,1]$. Workers solve tasks by applying knowledge of matching complexity. For any task of difficulty θ, output is strictly realized only if the worker's knowledge reserve z satisfies $z \geq \theta$. Without loss of generality, we assume that completing 1 unit measure of a task yields 1 unit of output. Consequently, the firm's total output is determined by the aggregate measure of completed tasks across the entire continuum.

2.2 Organizational Structure and Labor Allocation

The firm employs two types of labor endowed with heterogeneous knowledge reserves: rank-and-file employees (L) and executives (M), which together constitute a standard vertical hierarchy. For simplicity, the labor supply for both tiers is normalized to one. Task execution requires the mastery of corresponding knowledge. Due to individuals' bounded cognitive capacities and the increasing marginal time cost of knowledge acquisition, it is economically unfeasible for any single employee to master the entire knowledge spectrum. Therefore, the firm achieves efficient knowledge acquisition and problem-solving through a vertical division of labor, structured as follows:

Rank-and-file employees operate on the production frontline, initially encountering and executing tasks with complexity $\theta \in [0, z_L]$. These predominantly comprise routine tasks characterized by high frequency and standardization. When confronting exceptional difficulties—where task complexity exceeds their knowledge reserve $(\theta > z_L)$—rank-and-file workers cannot solve the problems independently. Instead, they escalate these tasks up the hierarchy, a process that incurs communication costs. Assuming the time endowment of rank-and-file employees is fully allocated to solving problems within their competence threshold, their total output equals z_L units.

Executives serve as the central coordinators of knowledge, responsible for resolving the higher-complexity tasks $\theta \in (z_L, 1]$ escalated by frontline workers. Although processing these escalated tasks consumes executive time, this vertical

division of labor substantially optimizes the utilization efficiency of scarce, advanced knowledge. The measure of the executives' exclusive task domain is defined as $v_M = 1 - z_L$. Given a unit time endowment and a communication cost parameter h (which detracts from available productive time), executive output is given by $(1 - z_L)(1 - h)$.

Total output is thus the linear aggregation of the outputs from both rank-and-file employees and executives. Consequently, the firm's total production function for a unit measure of tasks can be expressed as:

$$Y(z_L) = z_L + (1 - z_L)(1 - h) = 1 - h + hz_L \quad (1)$$

2.3 Knowledge Acquisition Costs and Compensation Setting

The production process relies on two distinct forms of knowledge: explicit and tacit. Explicit knowledge is codifiable and can be seamlessly transmitted through formalized, standardized channels (e.g., manuals, flowcharts, and program code) (Nonaka, 1994). Within the production process, it manifests as standard operating procedures. Conversely, tacit knowledge, as conceptualized by Polanyi (1966), encompasses strategic intuition, complex organizational coordination, unstructured problem-solving capabilities, and nuanced interpersonal judgment. Consequently, it is highly context-dependent, specialized, and inherently difficult to codify.

The acquisition of knowledge naturally incurs costs. Let $C_E(z_L)$ denote the learning cost function associated with explicit knowledge, with $c_e = \partial C_E / \partial z_L$ representing the marginal learning cost for specific tasks. Consistent with the knowledge hierarchy framework, the marginal cost of learning increases with task complexity, implying $\partial c_e / \partial z_L > 0$. The acquisition of tacit knowledge, however, transcends mere document assimilation, necessitating long-term learning-by-doing and profound cognitive investment.

Rank-and-file compensation (w_L): Rank-and-file employees rely predominantly on explicit knowledge to execute routine problems within the task allocation boundaries delineated by the firm. Their compensation structure is thus bipartite. The first component is the outside option value (reservation wage) in the broader labor market—

determined by exogenous factors such as labor market competition and institutional minimum wage floors—denoted by $\overline{w}$. The second component is the compensation C_E provided by the firm to subsidize employees' explicit knowledge acquisition. Accordingly, rank-and-file compensation is formulated as:

$$w_L = \overline{w} + C_E(z_L) \tag{2}$$

Since the firm lacks the economic incentive to subsidize learning costs beyond the requisite task scope, the maximum knowledge acquisition cost the firm is willing to cover for rank-and-file employees is strictly bounded by $C_E(z_L)$.

Executive compensation (w_M): The compensation-setting mechanism for executives diverges fundamentally from that of rank-and-file employees. Executives rely extensively on tacit knowledge to navigate exceptional, non-routine problems (e.g., mitigating sudden public relations crises, exercising strategic intuition, and orchestrating complex internal coordination). Given the highly context-dependent nature of tacit knowledge, it is inherently unverifiable; the firm cannot precisely measure the exact units of tacit knowledge an executive expends on an idiosyncratic problem.

Overcoming this unobservable pricing friction requires aligning incentives. According to residual claim theory, the firm must share its operating surplus with executives to properly incentivize the deployment of their tacit knowledge toward resolving these exceptional problems. Consequently, executive compensation is formulated as:

$$w_M = p + \gamma(v_M)[Y - w_L L - DK] \tag{3}$$

where p represents the outside option value in the executive labor market. The parameter $\gamma \in (0,1)$ captures the executive's bargaining power, which is strictly increasing in their span of control, such that $\gamma'(v_M) > 0$. The economic rationale is straightforward: the broader the task allocation boundaries controlled by executives, the greater the systemic impact of their deployed tacit knowledge and decision-making authority on aggregate firm production. This expanded control directly amplifies their bargaining leverage in the distribution of residual profits. In essence, managerial power and control over the production process directly translate into enhanced bargaining

power during compensation negotiations. Finally, K and D denote the firm's capital stock and the corresponding depreciation rate, respectively. Therefore, $[Y - w_L L - DK]$ constitutes the net operating surplus shared between the firm and its executives.

2.4 AI Technology and Firm Decision-making

We now incorporate AI technology into our theoretical framework. Within the knowledge-based hierarchical model, acquiring explicit knowledge traditionally necessitates substantial investments of time and cognitive effort through professional training. However, the generative capabilities of modern AI fundamentally alter the paradigm of knowledge retrieval, comprehension, and generation, substantially lowering the barriers to acquiring and deploying explicit knowledge. We therefore posit that AI adoption structurally reduces the explicit learning costs for rank-and-file employees. Let A denote the magnitude of the AI technological shock (i.e., the firm's level of AI application). The learning cost function for rank-and-file employees is thus reformulated as $C_E(z_L, A)$. The AI-induced reduction in absolute and marginal learning costs implies $\partial C_E/\partial A < 0$ and $\partial c_e/\partial A < 0$, respectively.

Furthermore, AI functions as a general-purpose technology characterized by capacities for self-evolution and iterative improvement (Baily et al., 2026). Deploying AI necessitates specialized physical carriers, denoted as fixed digital and intelligent capital equipment, $\overline{K_{AI}}$, subject to a depreciation rate $D(A)$. Given the self-evolutionary nature of AI, continuous technological iteration accelerates the obsolescence of corresponding capital assets, strictly implying $\partial D(A)/\partial A > 0$.

Given the exogenous AI technology level A, the representative firm seeks to maximize its total operating surplus, $S(z_L, A)$, by optimizing the task allocation boundary, z_L, between executives and rank-and-file employees. This optimization problem is formally expressed as:

$$max_{z_L}\ S(z_L, A) = [1 - h + hz_L] - w_L(z_L, A) - D(A)\overline{K}_{AI} - \Phi \qquad (4)$$

The first-order condition with respect to z_L is given by:

$$\frac{\partial S(z_L, A)}{\partial z_L} = h - c_e(z_L^*, A) = 0 \Longrightarrow h = c_e(z_L^*, A) \quad (5)$$

In equilibrium, the optimal task allocation boundary z_L^* between executives and rank-and-file employees is determined where the marginal cost of acquiring explicit knowledge exactly equals the marginal communication cost.

2.5 Asymmetric Impact of AI Technology on Firm Compensation

Deriving from the firm's surplus maximization conditions, we establish Proposition 1:

Proposition 1 (Downward Task Delegation Effect):

The optimal task boundary z_L^* shifts from executives to rank-and-file workers as AI adoption deepens ($dz_L^*/dA > 0$). AI lowers the cost of acquiring explicit knowledge, allowing frontline workers to take on tasks previously reserved for managers.

Proof. Applying the implicit function theorem to the first-order condition $h = c_e(z_L^*, A)$ with respect to A, holding communication costs constant, we obtain:

$$\frac{dz_L^*}{dA} = -\frac{\frac{\partial c_e}{\partial A}}{\frac{\partial c_e}{\partial z_L}} > 0 \quad (6)$$

The economic intuition underpinning Proposition 1 stems from AI's superior information processing and pattern recognition capabilities. This technological empowerment substantially lowers the cognitive threshold required for rank-and-file employees to master complex explicit knowledge (i.e., $\partial c_e/\partial A < 0$). Consequently, once traditionally complex analytical processes (e.g., financial modeling or generating market insights) are modularized and codified by AI, rank-and-file employees can acquire the requisite skills and execute these tasks at a significantly lower marginal cost. Nevertheless, consistent with our baseline assumptions, the marginal learning cost continues to rise with intrinsic task complexity (i.e., $\partial c_e/\partial z_L > 0$).

Proposition 2 (Non-Monotonic Dynamics of Total Operating Surplus):

Ceteris paribus, the firm's equilibrium total operating surplus is a non-monotonic function of the AI technology level A. AI adoption does not deterministically enhance operating surplus; the sign of dS^*/dA remains theoretically ambiguous.

Proof. By the envelope theorem, the derivative of the equilibrium total operating surplus with respect to A is given by:

$$\frac{dS^*}{dA} = -\frac{\partial C_E(z_L^*, A)}{\partial A} - \frac{\partial D(A)}{\partial A}\overline{K}_{AI} \tag{7}$$

Given $\partial C_E(z_L^*, A)/\partial A < 0$ and $\partial D(A)/\partial A > 0$, the net impact of AI adoption on the firm's surplus is indeterminate. If $-\partial C_E(z_L^*, A)/\partial A > [\partial D(A)/\partial A]\overline{K}_{AI}$, then $dS^*/dA > 0$. Here, the learning efficiency gains induced by AI dominate the depreciation costs, expanding the firm's operating surplus. Conversely, if $-\partial C_E(z_L^*, A)/\partial A < [\partial D(A)/\partial A]\overline{K}_{AI}$, then $dS^*/dA < 0$.

Because AI technology undergoes rapid and continuous advancement—evidenced by the exponential growth in computational demands and the rapid versioning of large language models—firms are thrust into a "Red Queen effect," akin to a technological arms race. To sustain continuous AI-driven efficiency gains, firms must persistently inject massive volumes of digital and intelligent capital. The obsolescence rate of this capital accelerates rapidly amidst ongoing technological iterations. When capital depreciation costs, catalyzed by this relentless iteration, outstrip the marginal efficiency gains, the firm's total operating surplus contracts. This contraction fundamentally precipitates a structural dissipation of rents.

Proposition 3 (Pay Gap Compression Effect):

The deployment of AI technology exerts an asymmetric impact on the compensation of executives and rank-and-file employees, resulting in a significant compression of the intra-firm pay gap. For rank-and-file employees, the optimal compensation contract provided by the firm is $w_L = \overline{w} + C_E(z_L^*, A)$; for executives, the compensation contract is $w_M = p + \gamma(1 - z_L^*)S(z_L^*, A)$. Let $\Omega = w_M/w_L$ denote the intra-firm pay gap between executives and rank-and-file employees. It follows that $\partial\Omega/\partial A < 0$.

Proof. Differentiating Ω with respect to the AI technology level A yields:

$$\frac{d\Omega}{dA} = \frac{1}{w_L^2}\left[\frac{dw_M}{dA}w_L - \frac{dw_L}{dA}w_M\right] \tag{8}$$

$$\frac{dw_L}{dA} = \frac{\partial C_E}{\partial z_L^*}\frac{dz_L^*}{dA} + \frac{\partial C_E}{\partial A} \approx 0 \tag{9}$$

$$\frac{dw_M}{dA} = S(z_L^*, A)\frac{\partial \gamma}{\partial z_L^*}\frac{dz_L^*}{dA} + \gamma(1 - z_L^*)\frac{dS^*}{dA} < 0 \tag{10}$$

For rank-and-file employees, the AI-induced reduction in learning costs exerts downward pressure on wages ($\partial C_E/\partial A < 0$). Conversely, the expansion of their task allocation boundaries and the heightened complexity of their responsibilities exert upward pressure on wages ($\frac{\partial C_E}{\partial z_L^*}\frac{dz_L^*}{dA} > 0$). Because rank-and-file compensation functions fundamentally as an implicit insurance contract, these upward and downward wage pressures largely offset one another in equilibrium, yielding $dw_L/dA \approx 0$.

For executives, given that bargaining power strictly increases with the span of control, a reduction in their exclusive task domain implies $\partial\gamma/\partial z_L^* < 0$. Consequently, the term $S(z_L^*, A)\frac{\partial \gamma}{\partial z_L^*}\frac{dz_L^*}{dA} < 0$ captures the downward effect on compensation driven by the compression of the executive span of control. Furthermore, $\gamma(1 - z_L^*)\frac{dS^*}{dA} < 0$ isolates the downward effect driven by the contraction of the total operating surplus. During the early-to-intermediate stages of AI capital investment and application, accelerated depreciation costs reliably outpace immediate efficiency gains, ensuring that $\gamma(1 - z_L^*)\frac{dS^*}{dA} < 0$ robustly holds. Therefore, $dw_M/dA < 0$.

Hypothesis 1 (Baseline Hypothesis):

AI shocks exert an asymmetric impact on compensation stickiness across the corporate hierarchy, thereby driving the compression of the intra-firm pay gap. Specifically, under AI shocks, executive compensation exhibits pronounced downward flexibility, whereas rank-and-file wages maintain stable downward rigidity. In essence, the traditional downward rigidity insulating executive compensation fractures in the AI era.

Hypothesis 2 (Mechanism Hypothesis):

AI shocks drive the downward flexibility of executive compensation primarily through the channel of positional rent dissipation. Theoretically, this rent dissipation is catalyzed by task boundary restructuring and the crowding-out effect of digital and intelligent capital. The specific mechanisms operate as follows:

First, AI's general-purpose predictive and information-processing capabilities substantially lower the threshold for acquiring explicit knowledge. This drives analytical and decision-making tasks—historically monopolized by executives—downward to the production frontlines. Consequently, the task allocation boundaries for rank-and-file employees expand, elevating the marginal value of their human capital. Coupled with the implicit insurance contracts inherent in rank-and-file labor relations, this expansion effectively offsets the substitution risks posed by AI automation for routine tasks, thereby solidifying the downward rigidity of rank-and-file compensation.

Second, this downward delegation of tasks and the subsequent flattening of information flows endogenously compress the executive span of control and erode their exclusive information advantages. This deconstruction of bureaucratic power directly diminishes the structural bargaining power executives hold over the distribution of the residual surplus. As a result, the positional rents traditionally derived from hierarchical advantages face marginal dissipation, stripping executive compensation of its institutional insulation against shocks. Furthermore, the acceleration of AI-driven capital deepening structurally crowds out labor's distributive share, substantially compressing the overall surplus pool available for distribution. Because the deep deployment of AI necessitates immense capital expenditures on computational infrastructure—which exhibits rapid endogenous depreciation—the burden of this surplus contraction is disproportionately borne by executives whose bargaining power over positional rents has already eroded. Ultimately, this rent dissipation manifests empirically as high downward flexibility in executive compensation.

3 Research Design

3.1 Policy Background

To empirically test the theoretical predictions regarding AI shocks and asymmetric compensation stickiness, the central empirical challenge lies in constructing an identification strategy capable of capturing exogenous AI technological shocks to micro-level firms, which subsequently induce heterogeneous adjustments in their internal distribution structures. Endogeneity is a first-order concern: a firm's AI adoption correlates with its operational conditions and existing digital infrastructure, so raw AI investment data do not identify causal effects. Given that the depth of a firm's AI application depends not only on internal strategic decisions but also critically on external digital infrastructure, algorithmic computing ecosystems, and the broader policy environment, we exploit the establishment of the National New Generation AI Innovation and Development Pilot Zones, spearheaded by the Ministry of Science and Technology (MOST), as an exogenous quasi-natural experiment representing AI shocks.

To accelerate the deep integration of AI with the real economy, MOST issued the Guidelines for the Construction of National New Generation AI Innovation and Development Pilot Zones (hereafter, the Guidelines) in August 2019. Between 2019 and 2021, 18 cities nationwide were successively approved to establish these pilot zones. Crucially, this policy extends beyond simple financial subsidies. By comprehensively developing frontier algorithmic platforms, computational infrastructure, and open-source data environments, it provisions the essential carriers for the continuous, iterative evolution of AI technologies among regional firms. This top-down national strategy, implemented with spatial and temporal variation, provides an ideal staggered quasi-natural experimental setting to evaluate the reshaping effect of AI technology on intra-firm hierarchical compensation structures.

According to the application criteria delineated in the Guidelines, candidate cities must satisfy stringent prerequisites, including abundant scientific and educational resources, robust infrastructure, and a core AI industry scale exceeding RMB 5 billion. Consequently, the digital technological endowments and economic capacities of the selected pilot cities are inherently superior to those of non-selected cities. This non-random, targeted policy deployment implies systematic differences in the initial

resource endowments and internal organizational structures between firms in the treatment and control groups, potentially leading to divergent evolutionary trajectories over time. If the endogenous evolutionary dynamics driven by this non-random site selection are not effectively isolated, the empirical model risks violating the parallel trends assumption.

3.2 Identification Strategy

To address the potential biases induced by staggered policy implementation and non-random site selection, we adopt a dual identification strategy. On the one hand, to alleviate concerns regarding arbitrary treatment-effect heterogeneity inherent in conventional two-way fixed effects (TWFE) regressions, we employ the robust and efficient imputation estimator proposed by Borusyak et al. (2024). This approach explicitly isolates the estimation of the potential outcome model using only untreated observations, thereby circumventing the contamination of treatment effects. On the other hand, to control for the non-randomness of the pilot zone selection, we extract the pre-policy sample means of key firm-level characteristics (including total assets, profitability, and property rights nature) and explicitly incorporate their interactions with linear time trends into the regression model. This strategy aims to absorb and isolate non-parallel evolutionary trajectories stemming from disparities in initial endowments, thereby precisely capturing the net causal effect of AI shocks on the intra-firm compensation distribution structure. Specifically, conditional on satisfying the parallel trends assumption, the implementation logic of this imputation estimator proceeds in three distinct steps:

First, estimate the model of untreated potential outcomes using untreated observations only. The estimation procedure relies exclusively on the sample unexposed to the policy intervention (i.e., untreated observations where $D_{it} = 0$, encompassing all periods for the control group and the pre-treatment periods for the treated group) to consistently estimate unit fixed effects, period fixed effects, and the time trends associated with baseline characteristics:

$$WageGap_{it}(0) = \alpha_i + \gamma_t + X'_{i,pre}\theta \times t + \varepsilon_{it} \qquad if \ D_{it} = 0 \qquad (11)$$

In Equation (11), $WageGap_{it}(0)$represents the untreated potential intra-firm pay gap for firm i in period t. α_i and γ_t denote unit (firm) and period fixed effects, respectively. To mitigate the endogeneity of site selection and avoid "bad control" bias from concurrent financial indicators, we extract the mean vector of baseline characteristics $X'_{i,pre}$ (comprising total assets, Tobin's Q, ROA, ownership concentration, asset-liability ratio, and property rights nature) prior to the comprehensive rollout of the policy. We then explicitly incorporate the interaction between this set of pre-determined variables and the time trend t to absorb the non-parallel evolutionary trajectories driven by differences in initial scale and resource endowments.

Second, impute the untreated potential outcomes and obtain an estimated treatment effect for each treated observation. Based on the unbiased estimates $(\hat{\alpha}_i, \hat{\gamma}_t, \hat{\theta})$ obtained in the first step, and leveraging the unit fixed effects $\hat{\alpha}_i$ identified from the pre-treatment periods of the treated group, we extrapolate the model to the treated observations ($D_{it} = 1$). This allows us to impute the counterfactual potential pay gap for each treated firm-year observation had the AI policy shock not occurred:

$$\widehat{WageGap}_{it}(0) = \hat{\alpha}_i + \hat{\gamma}_t + X'_{i,pre}\hat{\theta} \times t \qquad (12)$$

Third, estimate the target estimand by a weighted sum of the estimated treatment effects. Subtracting the imputed untreated potential outcome $\widehat{WageGap}_{it}(0)$ from the actual observed intra-firm pay gap $WageGap_{it}$ yields the micro-level heterogeneous treatment effect $\hat{\tau}_{it}$ for each treated firm-year observation. We average these estimates across all treated observations to obtain the average treatment effect on the treated (ATT), which is the target estimand of our empirical analysis:

$$\hat{\tau}_{ATT} = \frac{1}{N_{treat}} \sum_{D_{it}=1} \left(WageGap_{it} - \widehat{WageGap}_{it}(0) \right) \qquad (13)$$

where N_{treat} denotes the total number of treated observations in the sample.

To ensure valid and robust statistical inference, we cluster the standard errors of all regression coefficients at the prefecture-level city in our empirical analysis. The

primary rationale is twofold: First, regarding the systematic nature of policy implementation, the pilot zone policy is administered at the city level; consequently, firms within the same city inevitably face correlated unobservable regional concurrent shocks, such as variations in policy enforcement intensity and local financial subsidies. Second, regarding the structure of the error terms, clustering at the prefecture-city level effectively accounts for both serial correlation over time within firms and spatial correlation across firms within the same region, thereby preventing the severe underestimation of standard errors and ensuring the validity of our causal identification results.

3.3 Variable Selection

Dependent Variables: The primary dependent variables are the executive compensation level ($lnAMP$), the rank-and-file employee compensation level ($lnAEP$), and the intra-firm pay gap ($WageGap$). To empirically depict the intra-firm income distribution structure, we follow Kong et al. (2017) and measure the intra-firm pay gap ($WageGap$) as the ratio of average executive compensation (AMP) to average rank-and-file employee compensation (AEP). $WageGap$ measures compensation polarization across the firm's hierarchy — a direct manifestation of agency frictions in surplus distribution.

Specifically, average executive compensation (AMP) measures the average income of all senior executives, non-independent directors, and supervisors remunerated by the firm. It is calculated by dividing the "total annual compensation of directors, supervisors, and senior executives" by the "actual number of directors, supervisors, and senior executives receiving compensation" (explicitly excluding unpaid personnel and independent directors).

Conversely, average rank-and-file employee compensation (AEP) captures the average income of frontline employees, excluding executives. It is computed as the total compensation of all employees minus the total compensation of directors, supervisors, and senior executives, divided by the total number of employees minus the number of

directors, supervisors, and senior executives. Total employee compensation is rigorously defined as the cash paid to and on behalf of employees (extracted from the cash flow statement) plus the net change in the ending and beginning balances of employee benefits payable (extracted from the balance sheet).

Furthermore, to accurately isolate and examine the structural shock of AI technology on compensation levels across distinct vertical echelons, we apply the natural logarithm to these calculated metrics, denoting the executive compensation level as $lnAMP$ and the rank-and-file compensation level as $lnAEP$.

Independent Variable: AI shocks (AI). The core explanatory variable is a policy dummy indicating whether the prefecture-level city where the firm is registered has been approved as a National New Generation AI Innovation and Development Pilot Zone. Based on the phased rollout lists published by the Ministry of Science and Technology, we manually compiled a roster of approved cities. The assignment rule is as follows: if the city of the firm's registered address is approved to establish a pilot zone, the variable AI takes the value of 1 in the year of approval and all subsequent years; otherwise, it equals 0.

Given the non-sudden nature of national industrial policies, the establishment of these pilot zones entails comprehensive administrative approval procedures, encompassing preliminary planning, local applications, and field inspections. During this incubation and signal-release period, firms possess robust information acquisition and forward-looking decision-making capabilities, inherently inducing significant policy anticipation effects[1].

To prevent anticipatory behavior from contaminating the control group, we shift the treatment timing by one period in our baseline specification — a shift(1) design following Borusyak et al. (2024). This strategy explicitly incorporates anticipatory responses occurring one period prior to the shock into the treatment effect evaluation

[1] Appendix 1 reports the dynamic parallel trends test of the anticipation effect of the National New Generation AI Innovation and Development Pilot Zone policy on the intra-firm pay gap, alongside detailed theoretical support from the relevant econometrics literature.

window, thereby capturing the true net causal effect of the AI technological shock more accurately and without bias.

Control Variables.

As delineated in our identification strategy, to mitigate endogeneity concerns and effectively control for divergent evolutionary trajectories stemming from heterogeneous initial endowments, we compute the pre-policy sample means of key firm-level characteristics (2015–2017). These pre-determined baseline values serve as the foundation for constructing the time-trend interaction terms within our control vector[2].

Specifically, this pre-determined control vector comprises Firm Size ($Size$), Firm Value ($TobinQ$), Profitability (ROA), Capital Structure (Lev), Ownership Concentration ($OwnCon$), and State-Owned Enterprise status (SOE). Detailed variable definitions are provided in Table 1.

3.4 Data Sources and Descriptive Statistics

Our research sample consists of A-share listed firms on the Shanghai and Shenzhen stock exchanges. The sample period is strategically defined as spanning from 2015 to 2024. The selection of 2015 as the starting point aligns with the State Council's official issuance of the Guiding Opinions on Actively Promoting the "Internet Plus" Action in July 2015, which formally designated "Internet Plus AI" as a pivotal national strategic initiative.

Firm-level micro-data—encompassing executive compensation, rank-and-file compensation, and other pertinent financial and corporate governance variables—are primarily extracted from the China Stock Market & Accounting Research (CSMAR) database. The roster of approved National New Generation AI Innovation and

[2] Under standard specifications, the pre-treatment baseline period preceding the policy's comprehensive rollout would typically span from 2015 to 2018. However, having identified a significant policy anticipation effect and subsequently incorporated a one-period lead (shift(1)) for the treatment timing, firm characteristics in 2018 are highly susceptible to intervention contamination. To circumvent the "bad control" problem, we strictly restrict the extraction window for baseline characteristics to the uncontaminated period of 2015–2017. Consequently, all time-trend interactions for the pre-treatment characteristic variables in the subsequent empirical analyses utilize these 2015–2017 sample means.

Development Pilot Zones at the prefecture-city level was manually compiled and matched using official policy documents published sequentially by the Ministry of Science and Technology.

To ensure data quality and the robustness of our empirical estimates, we apply the following rigorous screening procedures to the initial sample: (1) excluding firms operating in the financial sector; (2) excluding firms designated as Special Treatment (ST or *ST) or possessing anomalous listing statuses during the observation period; (3) excluding technically insolvent firms (leverage ratio > 1), observations with obvious anomalies in the intra-firm pay gap (values < 1), and those missing core variables; (4) to ensure the valid identification of baseline unit fixed effects within the imputation estimator framework, excluding treated observations with an effective pre-treatment period of fewer than two years[3]; and (5) excluding firms with a continuous listing duration of fewer than two years over the entire sample period to guarantee the general identifiability of unit fixed effects. Finally, to mitigate the influence of extreme outliers on model estimation, all continuous variables are winsorized at the 1st and 99th percentiles. Following these procedures, we obtain a final unbalanced panel dataset comprising 28,723 valid firm-year observations across 3,161 listed firms[4].

Table 2 presents the descriptive statistics for the primary variables used in this paper. Panel A details the full sample distribution, while Panel B reports tests of differences in pre-treatment characteristics between the treatment and control groups, categorized by eventual exposure to the policy shock. As shown in Panel A, the core dependent variable—the intra-firm pay gap—has a mean of 4.842, a standard deviation of 3.461, and ranges from 1.207 to 21.574. This indicates a substantial disparity in the

[3] According to the underlying identification logic of the imputation estimator by Borusyak et al. (2024), the first step relies entirely on untreated observations to accurately fit unit and period fixed effects. Therefore, if the effective pre-treatment observation period for a treated unit is less than two years, its unit fixed effects cannot be identified (due to insufficient degrees of freedom), or the baseline estimation will suffer from excessive noise, precluding the inference of a reliable counterfactual state. Consequently, such observations must be excluded. In contrast, for pure control group units that are never exposed to the shock, the concept of a pre-treatment window is inapplicable. As long as their continuous duration spans at least two years, all their observations can be leveraged for the valid identification of aggregate time trends and thus do not require exclusion.

[4] Appendix 2 reports the structural characteristics of the unbalanced panel and the observation period distribution of the regression sample.

internal compensation distribution among China's listed firms, alongside considerable cross-sectional heterogeneity. Furthermore, the mean executive compensation level ($lnAMP$) is 13.193, whereas the rank-and-file compensation level ($lnAEP$) is 11.798, reflecting an overall distribution that aligns with standard economic intuition.

Panel B reveals that prior to the policy shock, the initial endowments of the treatment and control groups exhibited similar core profitability but divergent structural characteristics. Specifically, there is no statistically significant difference in baseline profitability (ROA) between the two groups, indicating that treated firms located in the pilot zones and control firms shared comparable initial profitability fundamentals. However, relative to the control group, treated firms generally exhibit larger asset scales, higher market valuations, moderately elevated leverage, and a significantly larger proportion of concentrated ownership and state-owned enterprise (SOE) backgrounds.

This distributional disparity in baseline characteristics underscores the necessity of uniformly incorporating the linear time trends of these pre-treatment variables into the imputation estimator model. On the one hand, controlling for the dynamic trends of variables such as firm size, capital structure, and property rights effectively absorbs systematic selection bias induced by differences in firms' initial structural endowments. On the other hand, although baseline profitability is balanced between the groups, profitability remains a core determinant of internal compensation contracts. Incorporating the time trend of pre-treatment profitability essentially constrains the treatment and control groups to consistent dynamic profitability expectation trajectories. This rigorous specification effectively controls for confounding interference arising from potential divergences in performance trajectories, ensuring that the estimated convergence of the intra-firm pay gap is unbiased and reliably captures the true causal effect.

4 Empirical Analysis Results

4.1 Baseline Regression Results

AI shocks compress the intra-firm pay gap. Table 3 reports the baseline estimates. Specifically, Columns (1) and (2) present the effects of AI shocks on the overall intra-firm pay gap, while Columns (3) and (4) disaggregate these effects to separately examine the compensation dynamics of executives and rank-and-file employees.

First, AI shocks significantly compress the intra-firm pay gap. Column (1) of Table 3 demonstrates that, after accounting for two-way fixed effects, the estimated coefficient of the core explanatory variable (AI) on the intra-firm pay gap is -0.318, which is statistically significant at the 1% level. To further mitigate potential endogeneity arising from the dynamic evolution of firms' ex-ante heterogeneous characteristics, Column (2) incorporates the linear time trends of the pre-treatment characteristic variables. The results indicate that the baseline estimate remains highly robust (-0.291, $p < 0.01$) upon the inclusion of these stringent controls. This provides robust empirical evidence supporting our primary hypothesis: AI shocks precipitate a structural compression of the intra-firm pay gap.

Second, the adjustments in compensation exhibit pronounced asymmetry across intra-firm hierarchical echelons. Columns (3) and (4) regress executive compensation levels ($lnAMP$) and rank-and-file compensation levels ($lnAEP$) on the AI shock, respectively. The estimates reveal a stark divergence. Column (3) indicates that AI shocks induce a significant reduction in executive compensation, with an estimated coefficient of -0.061 (significant at the 1% level). Conversely, Column (4) shows that the impact on rank-and-file compensation is a marginal -0.004 and completely lacks statistical significance.

These comparative findings provide compelling micro-level evidence for the emergence of asymmetric compensation stickiness under AI shocks. Confronted with the influx of digital and intelligent capital, rank-and-file compensation maintains its traditional downward stickiness, preserving its rigid baseline protections as an implicit insurance contract. In stark contrast, executive compensation—which inherently captures a substantial proportion of managerial rents and variable pay tied to the span of control—exhibits pronounced downward flexibility driven by rent dissipation. This implies that AI adoption does not necessarily impair the fundamental, underlying

productivity of the firm (as evidenced by the absence of pay cuts for frontline workers). Rather, by leveraging superior computational capabilities and substituting routine coordination tasks, AI effectively crowds out the excess bureaucratic positional rents historically entrenched at the apex of the corporate hierarchy. The subsequent mechanism analyses will empirically substantiate this rent dissipation channel.

4.2 Pre-trend Test and Dynamic Effects

An essential prerequisite for unbiased estimation in difference-in-differences (DiD) designs is that, prior to the AI shock, the untreated potential outcomes of the treatment and control groups exhibit parallel trajectories, thereby satisfying the parallel trends assumption. Having estimated the micro-level heterogeneous treatment effect $\hat{\tau}_{it}$ for each treated firm-year observation using the imputation estimator, we further apply an event-study design to test for parallel trends and evaluate the dynamic effects of AI shocks over time.

Unlike conventional two-way fixed effects (TWFE) DiD models that regress outcomes on a series of relative time dummies, dynamic effects under the imputation estimator are computed via a grouped weighted average of micro-level treatment effects across specific relative periods. Given that our baseline identification introduces a one-period lead to absorb policy anticipation effects, let E_i denote the year the city in which firm i is registered was officially approved as a pilot zone. The relative time for dynamic evaluation is defined as $k = t - (E_i - 1)$. Here, $k = 0$ corresponds to the period immediately preceding the official policy implementation (capturing anticipation), while $k \geq 1$ denotes the periods following the policy's official rollout. For our dynamic event window, we restrict the primary testing interval to span from five periods before to four periods after the policy intervention (i.e., $k \in [-4,4]$). The rationale for this restriction is to mitigate the biases associated with compositional changes at the tails of extended event windows in staggered DiD frameworks. Because there is a natural truncation between the staggered treatment years (2019–2021) and the sample boundary (2015–2024), distant relative periods beyond the $[-4,4]$ window

suffer from severe sample attrition and potential selection bias. Focusing on this core interval ensures the representativeness of the dynamic average treatment effects and effectively filters out estimation noise at the extremities of the unbalanced panel.

For any relative year k (encompassing pre-treatment periods $k \in \mathcal{K}_{pre}$ and post-treatment periods $k \in \mathcal{K}_{post}$), the average dynamic treatment effect $\widehat{\beta_k}$ is calculated as:

$$\hat{\beta}_k = \frac{1}{N_k} \sum_{(i,t)\in\Omega_k} \hat{\tau}_{it} \tag{14}$$

where Ω_k represents the set of treated observations at relative time k, N_k is the total number of observations within this set, and $\hat{\tau}_{it}$ is the individual treatment effect inferred by the imputation estimator.

Conventional pre-trend tests typically examine individual pre-treatment coefficients independently (i.e., $\beta_k = 0$ for $k \in \mathcal{K}_{pre}$). This approach, relying exclusively on pointwise significance, often suffers from insufficient statistical power, increasing the risk of failing to reject a violated parallel trends assumption. Furthermore, due to multicollinearity constraints, traditional dynamic DiD specifications must arbitrarily omit a base period, inherently limiting the completeness of the pre-trend evaluation. In contrast, the imputation estimator utilizes only untreated observations to fit the baseline fixed effects, enabling a consistent and complete estimation of placebo treatment effects for all pre-treatment relative periods. To rigorously verify the parallel trends assumption based on this framework, we target the dynamic effect parameter vector $\widehat{\boldsymbol{\beta}}_{pre}$ of the strictly pre-treatment periods (i.e., historical periods strictly preceding the anticipation phase) to construct the following joint null hypothesis: $H_0: \beta_{-4} = \beta_{-3} = \beta_{-2} = \beta_{-1} = 0$.

Following the robust testing procedure proposed by Borusyak et al. (2024), we construct the Wald statistic for this joint significance test:

$$W = \widehat{\boldsymbol{\beta}}'_{pre} \left(\widehat{\mathbf{V}}_{pre}\right)^{-1} \widehat{\boldsymbol{\beta}}_{pre} \tag{15}$$

where $\widehat{\boldsymbol{\beta}}_{pre}$ is a column vector comprising the estimated coefficients for each pre-treatment period, and $\widehat{\mathbf{V}}_{pre}$ is the estimated cluster-robust covariance matrix of these coefficients. If the p-value of this Wald test exceeds the conventional significance threshold, the null hypothesis cannot be rejected, statistically indicating that the sample strictly satisfies the parallel trends assumption.

The dynamically estimated coefficients for the intra-firm pay gap, along with their 95% confidence intervals, are presented in Figure 1. As the figure illustrates, during the strictly pre-treatment periods (relative years $k \in [-4, -1]$), regardless of whether the time-trend interactions of pre-treatment characteristics are included, the estimated coefficients for the intra-firm pay gap hover near zero, with all confidence intervals robustly containing zero. Crucially, the Wald joint significance tests yield p-values of 0.419 (Figure 1, left) and 0.321 (Figure 1, right). Both values comfortably exceed the 10% significance level, meaning the null hypothesis—that the pre-treatment dynamic coefficients are jointly zero—cannot be rejected. This evidence confirms that, after accounting for multi-way fixed effects and the linear time trends of pre-treatment characteristics, there are no significant differential time trends between the treatment and control groups prior to the onset of policy anticipation. Thus, the quasi-natural experimental design satisfies the fundamental premise of the parallel trends assumption.

Conditioned on the verification of parallel trends, the policy effects exhibit distinct dynamic evolutionary characteristics. Notably, because our baseline identification incorporates a one-period lead, $k = 0$ effectively captures the anticipatory response prior to the official implementation of the policy. Observing Figure 1, the intra-firm pay gap exhibits a significant negative response as early as period $k = 0$. This suggests that upon detecting the policy signal regarding the AI pilot zone approval, firms rapidly initiated forward-looking adjustments to their compensation structures. As the policy advances ($k \geq 1$) and AI technology is deeply integrated into operations, this compression effect demonstrates robust temporal persistence, maintaining a steady downward trajectory over the subsequent one to four years. This dynamic pattern indicates that the reshaping of intra-firm compensation structures by AI technology is

not a transitory impulse, but rather a structural transformation underpinned by long-term mechanisms.

4.3 Structured Placebo Tests

Were the baseline results driven by omitted variables? We conduct placebo tests following Eggers et al. (2024) across three dimensions. In the applied literature utilizing difference-in-differences identification strategies, robustness checks often rely on routine sensitivity tests, such as substituting core variables or restricting the sample period. However, to rigorously assess the plausibility of the core identification assumptions, it is necessary to exclude the interference of unobservable contemporaneous macroeconomic shocks or time-varying confounding factors at the firm level. Drawing upon the formal placebo test framework proposed by Eggers et al. (2024), this paper systematically re-evaluates the baseline findings across three methodological dimensions: placebo outcomes, placebo treatments, and placebo populations.

4.3.1 Placebo Outcome Test

A placebo outcome test involves re-estimating the baseline model by substituting the true outcome variable with a placebo outcome. This requires selecting an indicator that theoretically should not be affected by the treatment (AI shocks), thereby satisfying the "No Average Treatment Effect" (NATE) assumption. By anticipating statistically insignificant results, this test assesses whether the baseline model has inadvertently captured spurious structural changes, thereby alleviating concerns about spurious causal correlations arising from potential omitted variable bias or unobservable systematic trends. Following this logic, we select indicators that are theoretically orthogonal to AI shocks but highly susceptible to the same potential unobserved confounders as our primary dependent variable. Specifically, the placebo outcome test is conducted across two dimensions: firms' natural evolutionary characteristics and macro-level universal policy shocks.

First, we rule out the interference of firms' natural evolution and fluctuations in the external governance environment. We select the average age of executives and the number of corporate shareholders as placebo outcomes. Theoretically, the demographic aging of the executive cohort and the distribution of external ownership structures possess strong natural evolutionary attributes and should not be directly affected by short-term AI shocks. The regression results in Columns (1) and (2) of Table 4 demonstrate that the estimated coefficients of the AI shock on these two placebo indicators are highly attenuated and statistically insignificant. This result effectively rules out the possibility that the baseline model is merely capturing the natural evolution of the firm life cycle or cyclical fluctuations in the external capital market.

Second, we rule out the interference of overlapping macro-level universal policies. The implementation of regional industrial policies is frequently accompanied by compound policy interventions, such as universal fiscal subsidies or tax incentives. To test whether the National New Generation AI Innovation and Development Pilot Zone policy inadvertently exerted a direct resource compensation function, we select the effective tax rate and the ratio of non-operating income as additional placebo outcomes. The results in Columns (3) and (4) of Table 4 indicate that the AI policy shock has no significant impact on these two financial indicators. This finding effectively isolates the interference of policy spillover effects, such as fiscal support or tax relief, further corroborating that the compression of the intra-firm pay gap genuinely stems from the reshaping of the compensation structure driven by technological change, rather than a distributional effect triggered by the injection of exogenous fiscal resources.

4.3.2 Placebo Treatment Test

The placebo treatment test evaluates the robustness of causal identification by artificially constructing pseudo-policy shocks across temporal or spatial dimensions. By scrambling the true policy allocation mechanism, we generate an empirical distribution of estimated coefficients under pseudo-treatment states. This allows us to assess whether the significant effects observed in the baseline regression are statistically distinct from randomly generated noise, thereby alleviating concerns that the baseline

conclusions might be driven by unobservable time-varying confounding factors at the city level, pure statistical chance, or inherent biases in the model specification.

Specifically, while strictly maintaining the true time-frequency distribution and cross-sectional scale of the National New Generation AI Innovation and Development Pilot Zone policy, we randomly reshuffle the true policy implementation years for each city within the sample to construct spatiotemporally misaligned pseudo-treatments. Subsequently, utilizing the same imputation estimator as the baseline model, we conduct 500 Monte Carlo simulation regressions based on this randomized assignment process.

Figure 2 plots the kernel density distribution of the pseudo-treatment effect coefficients generated from these 500 random iterations. As depicted, the empirical distribution of the simulated coefficients is strictly centered around zero, exhibiting the characteristics of a standard normal distribution. This confirms that under randomly assigned counterfactual scenarios, no significant treatment effects emerge. More crucially, the true baseline regression coefficient (denoted by the vertical dashed line at -0.291) falls squarely in the extreme left tail of this simulated distribution. This provides robust distributional evidence that the net causal effect identified in this paper—where AI shocks significantly compress the intra-firm pay gap—is a non-random event, demonstrating strong causal validity.

4.3.3 Placebo Population Test

The placebo population test involves constructing subsamples that are theoretically immune to the core explanatory variable (AI shocks) and re-estimating the baseline model on these specific populations. By verifying the absence of the treatment effect within these immune subsamples, we can inversely confirm that the compression of the intra-firm pay gap observed in the baseline regression is indeed driven by the penetration of AI technology. This strategy mitigates concerns that the baseline findings might merely capture macroeconomic downturns, universal labor market fluctuations (such as systemic wage compression), or other unobservable macroeconomic common trends, thereby providing exclusive evidence for the technology-driven nature of the

effect. Following this logic, we precisely construct three types of placebo subpopulations—based on both institutional and technological dimensions—that are theoretically insulated from the direct impact of AI policies. The estimation results are reported in Table 5.

First, we construct a state-owned enterprise (SOE) subsample, which is institutionally immune. Constrained by strict administrative regulations, such as executive compensation caps within the state system, the compensation contracts of SOE executives fundamentally lack the capacity for expropriating market-oriented excess rents. Consequently, the task restructuring and rent dissipation effects induced by AI technology lack the requisite micro-foundations to materialize within state-owned enterprises. We therefore expect the effect of AI shocks on the SOE sample to be statistically insignificant. Segregating the sample based on ultimate controller status, Column (1) of Table 5 demonstrates that the estimated coefficient of the AI shock is entirely insignificant within the SOE subsample, validating this theoretical conjecture.

Second, we construct specific traditional-industry subsamples that are technologically immune. We select six major traditional sectors that rely heavily on physical operational environments and intensive interpersonal interactions (agriculture, forestry, animal husbandry, and fishery; accommodation and catering; healthcare; culture, sports, and entertainment; production and supply of electricity, heat, gas, and water; and public facilities and environmental management; corresponding to industry codes A, H, Q, R, D, and N). Because the production and management processes within these industries lack digitized scenarios amenable to AI algorithms, substantial natural barriers impede the deployment and penetration of AI technology, effectively severing the transmission channel of the technological shock at its source. The results in Column (2) of Table 5 confirm that, within this technologically immune population of specific traditional industries, the coefficient of the core explanatory variable remains statistically insignificant.

Third, we construct an industry subsample based on low technological dependence. Utilizing data from the 2015 National Input-Output Table, we calculate the direct consumption coefficients of all national economic industries for two core information

technology sectors: the manufacturing of communication equipment, computers, and other electronic equipment; and the information transmission, software, and information technology services industry. We then isolate the industries situated in the bottom 20% of these coefficients across the full sample. The production functions of these industries exhibit minimal intermediate consumption of information technology and rarely incorporate complex hardware/software procurement or digital coordination into daily operations. This renders them an immune population that is only marginally exposed to AI shocks. As anticipated, the results in Column (3) of Table 5 demonstrate that the estimated coefficient of the AI shock on the intra-firm pay gap is minuscule and fails to achieve statistical significance.

4.4 Other Robustness Checks

4.4.1 Excluding the Interference of Contemporaneous Policies

Excluding the interference of the State-Owned Enterprise (SOE) compensation cap policy. The commencement of our sample period in 2015 coincides with the comprehensive implementation of the Reform Plan for the Compensation System of Executives in Centrally Administered Enterprises (i.e., the SOE executive compensation cap mandate). This policy exogenously curtailed the excess compensation of SOE executives, thereby directly impacting the convergence of the intra-firm pay gap. Although the baseline regression controls for the linear time trends of pre-treatment property rights characteristics, a simple linear trend may insufficiently absorb the non-linear dynamic interference of this targeted mandate. To address this, we adopt two more rigorous identification strategies: First, because the compensation cap specifically targets state-owned entities, we exclude all SOEs and re-estimate the model utilizing a subsample of non-SOEs entirely unaffected by this policy constraint (Table 6, Column 1). Second, we introduce joint fixed effects of property rights nature and year into the full-sample model to absorb heterogeneous macroeconomic policy shocks faced by firms with differing property rights attributes across time (Table 6, Column 2). The results in Table 6 substantiate the robustness of the baseline findings.

Whether restricting the sample to non-SOEs or controlling for high-dimensional property rights-by-year interaction effects, the convergence effect of AI shocks on the intra-firm pay gap remains significant at the 1% level, with no substantial attenuation in the coefficient estimates.

Excluding the interference of the National AI Innovative Application Leading Zone and National Big Data Comprehensive Pilot Zone policies.

During our observation period, China concurrently approved and established National AI Innovative Application Leading Zones and National Big Data Comprehensive Pilot Zones. To prevent these parallel, multi-period overlapping policies from introducing confounding bias or cross-contamination into our treatment effects, we directly incorporate time-varying treatment status indicators for these two contemporaneous policies into the initial stage of the imputation estimator. This specification enables the imputation estimator to leverage "clean" observations—those not yet exposed to the AI Innovation and Development Pilot Zone shock—to unbiasedly isolate and identify the independent effects of these overlapping policies. Subsequently, during the counterfactual imputation of $Y(0)$ in the second stage, the compensation fluctuations induced by these contemporaneous policies are effectively purged. The results in Columns (3) and (4) of Table 6 demonstrate that after filtering out the interference of these concurrent multi-period DiD policies, the estimated coefficient of our core explanatory variable remains highly significant and robust.

4.4.2 Excluding the Interference of Firm-Specific Demand Shocks

A potential threat to our baseline identification is that the observed decline in executive compensation might merely reflect passive downward compensation adjustments triggered by coincidental macroeconomic downturns or firm-specific demand contractions, rather than structural rent dissipation. To isolate the AI technological shock from performance-driven interference, we anchor the time window prior to the release of the policy signal (2015–2017) and calculate the pre-treatment average operating revenue growth rate for each firm. Based on this metric, we filter the dataset to retain only a subsample of firms exhibiting ex-ante demand expansion (i.e.,

mean operating revenue growth rate > 0). As reported in Column (1) of Table 7, after excluding firms experiencing demand contraction, the policy estimation coefficient remains significantly negative at the 1% level. Even among firms with strong pre-treatment growth — those under no demand-contraction pressure — AI shocks significantly narrow the pay gap. This rules out the alternative that passive performance-induced pay cuts drive our results.

4.4.3 Controlling for Time-Varying Regional Unobservables

Given that the site selection for the national pilot zones is not entirely random and frequently favors regions with superior digital infrastructure and robust economic fundamentals, the dynamic evolution of these macro-level characteristics might simultaneously influence intra-firm compensation decisions. To address this, we augment the baseline model—which already includes pre-treatment firm characteristics—by introducing pre-treatment macro-characteristic variables at the prefecture-city level. These encompass the regional economic development level (the natural logarithm of per capita GDP), industrial structure (the ratio of secondary to tertiary industry value-added), internet broadband penetration (the ratio of internet users to the registered population), and government fiscal capacity (the ratio of general public budget expenditure to regional GDP). We explicitly incorporate the interactions between these regional variables and linear time trends into the control vector. As shown in Column (2) of Table 7, after absorbing these systematic regional evolutionary trends, the baseline conclusion remains entirely robust to the interference of city-level, time-varying unobservable factors.

4.4.4 Excluding Nominal Inflation Effects

Given that our panel spans a 10-year observation window, the aggregate nominal compensation and financial metrics of firms are inevitably subject to the interference of macroeconomic price inflation, which could precipitate measurement bias in the related variables. To effectively strip away the inflation effect embedded in nominal

prices, we anchor the first year of the sample period (2015) as the baseline and utilize the national annual Consumer Price Index (CPI) to construct a fixed-base CPI series. We then use this index to deflate all continuous nominal financial and compensation variables into real values. The results in Column (3) of Table 7 confirm that under a pure real-value measurement—fully purged of inflationary distortions—the convergence effect of AI shocks on the intra-firm pay gap remains robust.

4.4.5 Mitigating Spatial Self-Selection Bias

In evaluating regional industrial policies, a common empirical concern is that firms might strategically relocate their registered addresses to policy-targeted cities to extract fiscal subsidies and institutional dividends. Such strategic relocation constitutes a severe spatial self-selection bias. To ensure the quasi-exogeneity of the policy shock, we rigorously audit the records of changes in firms' registered addresses during the sample period. We subsequently exclude all firm observations that underwent cross-prefecture-city address changes between 2015 and 2024. As demonstrated in Column (4) of Table 7, after excluding these spatially relocated firms, the core estimation results undergo no substantial changes, effectively ruling out the interference of firm spatial sorting effects.

4.4.6 Addressing Compositional Bias

The non-random entry and exit of micro-level firms in panel data can trigger compositional effects. This implies that the convergence of the intra-firm pay gap observed in the baseline model might not stem from a genuine, technology-driven reshaping of intra-firm behavior, but rather from a statistical artifact caused by the systematic attrition and replacement of firms with specific compensation distribution profiles. While our baseline regression utilizes an unbalanced panel to maximize informational retention and accurately reflect capital market dynamics, we implement a stringent restriction to further validate our findings. We limit the sample to a strongly balanced panel, retaining only firms that existed continuously throughout the entire

observation period (2015–2024). The results in Column (5) of Table 7 show that after eliminating the potential compositional interference of firm entry and exit, although the number of valid observations drops to 19,640, the estimated coefficient of the AI shock remains remarkably stable, further solidifying the robustness of our core conclusion.

4.4.7 Excluding the Interference of Major Public Health Emergencies

Because our sample observation period encompasses the major global public health emergency of the COVID-19 pandemic, and the intensive approval period for the AI Pilot Zones highly overlaps with the pandemic in the time series, the baseline estimation is susceptible to confounding interference. The exogenous shock of the pandemic could plunge certain firms into operational distress, thereby triggering passive downward adjustments to executive compensation that might be statistically misidentified as the causal effect of AI shocks. To isolate our findings from the interference of this major emergency, we conduct robustness checks across both spatial and temporal dimensions; the results are reported in Table 8.

In the spatial dimension, we employ two distinct strategies. First, following the approach of Wang and Tang (2024), we completely exclude firm samples located in high-risk cities. Second, we introduce the pre-pandemic (2019) geographic distance of the firm's city to Wuhan, alongside the city's population density, as cross-sectional characteristics, interacting them with linear time trends in the control vector. The results in Columns (1) and (2) of Table 8 demonstrate that after absorbing these spatially heterogeneous evolutionary trends in pandemic risk exposure, the core estimated coefficient remains robust.

In the temporal dimension, we conduct a time-truncation test on the policy effect. If the convergence of the intra-firm pay gap observed in the full sample were primarily driven by passive downward compensation adjustments triggered by the public health emergency, the policy effect should theoretically only manifest following the outbreak of the pandemic. To test this, we partition the treatment effect into two distinct phases: pre-pandemic (τ_{pre}, before 2020) and post-pandemic (τ_{post}, 2020 and beyond). The results in Column (3) of Table 8 indicate that AI shocks exerted a significant

convergence effect on the intra-firm pay gap in both periods. More crucially, the coefficient of the inter-group difference between these two phases ($\tau_{post} - \tau_{pre}$) is completely statistically insignificant. This compellingly confirms that the policy effect maintained a high degree of temporal continuity, verifying that the observed compression of the intra-firm pay gap is a persistent structural shift rather than a short-term statistical artifact induced by the pandemic shock.

5 Mechanism Analysis

Baseline results show that AI shocks narrow the pay gap. We now ask: through what channels? We test the rent dissipation mechanism using firm-level data.

5.1 Direct Evidence from Rent Dissipation Mediators

Our theoretical framework posits that by reshaping task boundaries and fostering capital deepening, AI shocks trigger organizational flattening and non-linear shifts in surplus space, thereby dissipating managerial positional rents. To directly test this rent dissipation channel, we employ executive excess compensation (explicit rents) and excess perquisites (implicit rents) captured during the actual distribution process as direct proxies.

Regarding variable measurement, we aim to isolate the pure rent component by purging the effects of standard economic determinants. For explicit excess compensation, following Core et al. (1999) and Quan et al. (2010), we construct an executive compensation expectation model. By regressing out the influence of objective fundamentals—such as firm size, operating performance, leverage ratio, and industry environment—we extract the residual term and define it as executive excess compensation ($ExcessPay$).

For implicit perquisites, management may extract hidden private benefits through redundant, non-business-related expenditures. Drawing on the methodology of Chu and Fang (2016), we utilize an administrative expense stripping approach. After deducting compensation and amortization items and standardizing the data, we introduce firm

characteristics—such as asset structure, revenue volatility, and employee scale—into the expectation model. The resulting regression residual serves as our measure of executive excess perquisites ($ExcessPerks$). A significant contraction in these two excess return indicators would directly verify that AI shocks substantially compress executive rent-seeking space and positional rents.

Table 9 reports the test results for the rent dissipation mediators. First, we examine the curtailment effect of AI shocks on micro-level rents. Column (1) shows that the coefficient on the AI shock regarding executive excess compensation is -0.046, significant at the 1% level. Column (2) indicates that the coefficient on AI technology concerning executive excess perquisites is -0.008, also significant at the 1% level. These findings suggest that rather than indiscriminately reducing all managerial remuneration, AI technology specifically squeezes executives' excess returns, thereby confirming rent dissipation at the micro-phenomenological level.

Second, we verify the transmission pathway from these entity-level rents to the convergence of the pay gap. To mitigate omitted variable bias, we examine the partial correlation between the rent proxies and the pay gap, controlling for time-varying fundamental characteristics and two-way fixed effects. The results in Columns (3) and (4) of Table 9 demonstrate that the regression coefficients of executive excess compensation and excess perquisites on the intra-firm pay gap are 4.539 and 0.519, respectively, both significant at the 5% level or higher. This empirical evidence further substantiates the micro-foundation for the reshaping of distribution contracts: the excess returns extracted by management essentially serve as a critical pillar for maintaining their asymmetric compensation premiums and the internal distribution divide.

Synthesizing the empirical findings in Table 9, the underlying logic of our study is effectively validated. Specifically, AI shocks significantly suppress management's explicit excess compensation and implicit perquisites. The contraction of these excess returns objectively erodes the resource foundation upon which management relies to maintain downward compensation rigidity, which consequently drives the convergence of the intra-firm pay gap. This not only provides direct empirical support for our core

hypotheses but also highlights, from the perspective of primary distribution, the instrumental value of AI technology in mitigating intra-firm agency conflicts and optimizing the primary income distribution structure.

5.2 Suggestive Evidence Based on Financial and Organizational Structure

The preceding mediation regressions confirm that AI shocks trigger the substantial dissipation of managerial positional rents (i.e., excess compensation and excess perquisites). To further explore the underlying drivers behind this reshaping of distribution dynamics, this section provides suggestive evidence regarding adjustments to the firm's fundamental financial boundaries and organizational structures. Specifically, it aims to reveal how AI technology compresses executives' excess returns by tightening financial and organizational constraints across two main dimensions.

First, capital deepening and tightened financial constraints squeeze the resource boundaries of positional rents. Theoretically, a firm's extensive integration of AI technology is often accompanied by massive capital expenditures in digital and smart assets, such as hardware and software procurement, as well as algorithmic model training. Consistent with the free cash flow hypothesis, as technological iteration accelerates, the high depreciation and amortization associated with such capital investments inevitably erode the total operating surplus available for current-period distribution. This contraction of surplus space imposes a hard financial constraint, substantially weakening the financial foundation upon which management extracts excess remuneration via bureaucratic power. The results in Column (1) of Table 10 support this logic: the estimated coefficient on AI shocks for depreciation and amortization ($DepAmort$) is 0.049, significant at the 10% level. This indicates that the capital deepening induced by AI technology significantly compresses the firm's current distributable surplus, effectively establishing an objective financial boundary constraint that drives the dissipation of positional rents.

Second, the downward delegation of tasks and subsequent organizational restructuring dismantle the foundation for traditional hierarchical positional rents. The downward rigidity of managerial compensation largely stems from the bargaining power executives derive from their hierarchical advantages over specific complex tasks and information nodes. Because AI technology lowers the barriers to acquiring explicit knowledge and reduces information processing costs, firms are increasingly inclined to delegate easily learnable complex tasks to rank-and-file employees. This downward shift in task allocation not only streamlines redundant managerial nodes but also significantly mitigates the physical and informational frictions associated with cross-hierarchical communication. To empirically verify this conjecture, we construct proxies for the operational frictions within the bureaucratic organization. Specifically, the ratio of deputy executives ($DeputyRatio$) measures the proportion of auxiliary managerial nodes within the top management team, while the office administrative expense ratio ($OfficeExpRatio$) and the travel administrative expense ratio ($TravelExpRatio$)—calculated using detailed footnotes from financial statements—serve as additional proxies.

The results in Columns (2) through (4) of Table 10 demonstrate that AI shocks significantly promote organizational flattening and reduce internal coordination frictions. The coefficient on AI technology regarding the deputy executive ratio is -0.824 (significant at the 5% level); the coefficients for the office and travel administrative expense ratios are -0.033 and -0.021, respectively (both significant at the 1% level). These findings suggest that AI shocks not only substantively eliminate auxiliary management positions but also simultaneously compress the daily operational frictions inherent in bureaucratic systems. Thus, by driving structural flattening, technological shocks provide further corroboration for the deductions of our theoretical model: AI technology shrinks the exclusive task boundaries of management, directly eroding their compensation bargaining power derived from broad task spans.

Synthesizing the suggestive evidence in Table 10, we clarify the fundamental logic underlying the rent dissipation mechanism. AI technology's reshaping of the internal distribution contract is deeply rooted in shifts within the firm's underlying resources

and organizational structures, manifested in two key ways. First, elevated depreciation and amortization expenses squeeze management's surplus-sharing space via tightened financial constraints. Second, the reduction of deputy positions and the decline in friction costs erode management's positional advantages and bargaining power via organizational constraints. It is precisely the simultaneous contraction of these resource and organizational boundaries that strips management of the material and institutional foundations necessary to maintain excess compensation and perquisites, thereby cross-validating the direct evidence presented earlier.

5.3 Heterogeneity Tests Based on Constraints

The preceding analysis primarily establishes the convergence effect of AI shocks on the intra-firm pay gap at the aggregate level. To further demonstrate that this effect operates specifically through the rent dissipation mechanism, this section exploits cross-sectional heterogeneity. If rent dissipation drives our results, the compression effect should be stronger where agency conflicts are severer, outside options weaker, coordination pressure higher, or organizational resistance lower.

Regarding the estimation strategy, rather than adopting the traditional approach of splitting the sample for simple subsample regressions, we utilize the imputation estimator within a full-sample framework to conduct joint heterogeneity tests. Following Borusyak et al. (2024), we map this approach into triple-differences designs. In this framework, the data possess a second cross-sectional dimension—the heterogeneity group $g(i)$—in addition to the firm i and period t. Correspondingly, we strictly control for the interaction terms between grouping variables and periods (equivalent to the $\alpha_{g(i)t}$ component in the model of untreated potential outcomes). This design effectively avoids the negative weighting problem inherent in traditional staggered two-way fixed-effect (TWFE) regressions, prevents counterfactual baseline

fragmentation, and rigorously ensures the parallel trends assumption holds across all sub-dimensions[5].

First, we examine the heterogeneity driven by internal information asymmetry. According to our theoretical framework, positional rents heavily depend on the excess operating surplus extracted by management leveraging their insider control and information asymmetry. If this mechanism holds, the corporate governance value of AI technology should be directly proportional to its capacity to mitigate information asymmetry. Therefore, in firms with more opaque information environments, management's initial hidden rent-seeking space is larger, and the marginal compression effect induced by AI technology should be more substantial. We verify this inference utilizing the firm's internal information environment, a crucial determinant of principal-agent costs. Compared to firms with highly transparent information, those with poor disclosure quality suffer from pronounced internal and external information barriers. This amplifies the information asymmetry between shareholders and management, facilitating executives' exploitation of informational advantages to extract excess private benefits. However, AI's underlying data penetration and information integration capabilities can precisely and effectively bridge this gap. Consequently, we hypothesize that the convergence effect of AI shocks on the intra-firm pay gap is more pronounced in firms with poorer internal information environments. We measure this using the information disclosure evaluation ratings of listed companies on the Shenzhen Stock Exchange (1 = Excellent; 2 = Good; 3 = Qualified; 4 = Unqualified). Firms with a pre-treatment average rating leaning toward qualified or unqualified (mean > 2) are classified into the high information asymmetry group, and the rest into the low information asymmetry group. The results in Column (1) of Table 11 show that in the

[5] Utilizing the imputation estimator to conduct joint heterogeneity tests within a triple-differences design offers three primary econometric advantages: First, it provides a unified counterfactual baseline and higher statistical power. By jointly estimating fixed effects using all never-treated and not-yet-treated observations within the full sample, it avoids the loss of degrees of freedom typical of direct sample-splitting. Second, it avoids negative weighting. Unbiased heterogeneous treatment effects are obtained by appropriately averaging the observation-specific causal effect estimates derived from the imputed untreated potential outcomes, thereby bypassing the forbidden comparisons where late-treated units are inappropriately compared to already-treated reference groups. Third, it provides a strict defense of parallel trends. By explicitly incorporating the interactions between grouping variables and continuous year trends ($\alpha_{g(i)t}$), the model effectively absorbs group-specific time trends.

sample with poorer information environments, AI shocks significantly reduce the intra-firm pay gap, whereas this effect is relatively muted in firms with better information environments. The significant inter-group coefficient difference provides corroborative evidence for the positional rent dissipation mechanism.

Second, we explore the heterogeneity of bargaining power driven by outside options. Management's outside options in the managerial labor market are a vital source of leverage in compensation negotiations. Based on resource dependence theory, executives holding concurrent directorships in other listed companies possess substantial social capital, which substantially bolsters their bargaining power in internal distribution games by providing abundant outside options. If AI shocks indeed dissipate task rents by restructuring routine internal management functions, then executives lacking outside options—whose bargaining power is highly tethered to internal management tasks—should suffer a more severe compensation shock. Conversely, executives with strong outside options can leverage their external scarcity to build bargaining barriers, thereby exhibiting stronger compensation resilience. To verify this, we measure external network size by the average number of concurrent directorships held by the executive team in other listed companies, dividing the sample into high and low groups based on the pre-treatment median. The regression results in Column (2) of Table 11 demonstrate that for executive teams with weaker outside options, the convergence effect of AI shocks on the intra-firm pay gap is stronger, with a significant inter-group difference compared to the high outside option cohort. This indicates that the compensation of executives with abundant external resources exhibits strong non-substitutability, inversely corroborating that AI shocks precisely curtail the rents attached to firm-specific internal tasks.

Third, we test the heterogeneity of organizational coordination pressure based on factor endowment structures. The effectiveness of AI-driven organizational flattening may depend on the firm's initial organizational coordination pressure. Under the knowledge hierarchy framework, organizational layers exist to resolve exceptional problems escalated by frontline employees. If AI shocks reduce reliance on middle management through the downward delegation of tasks, the marginal optimization

effect of AI on the intra-firm pay gap should be more significant in organizations with higher initial coordination frictions and greater bureaucratic communication burdens. Generally, a firm's factor endowment structure dictates its information flow processing mode and coordination pressure. Compared to capital-intensive firms, the daily operations of labor-intensive firms rely heavily on the vertical coordination and hierarchical communication of a large workforce, inherently resulting in high initial coordination pressure. By algorithmizing explicit knowledge previously siloed at the management level, AI lowers cognitive barriers, enabling frontline employees to autonomously resolve exceptional difficulties that previously required upward escalation. This implies that for labor-intensive firms facing immense initial coordination pressure, AI technology unleashes a relatively stronger deconstruction effect, achieving deep organizational flattening and task rent dissipation through the large-scale streamlining of middle management nodes. Following Zhao et al. (2023), we use the capital-labor ratio (net fixed assets divided by total employees) to measure capital intensity, partitioning firms into high and low groups based on the pre-treatment median. The results in Column (3) of Table 11 show that in industries with lower capital intensity, AI shocks significantly contract the intra-firm pay gap, and this effect is significantly larger than in the high capital intensity group. This suggests that AI's substitution effect on management tasks has broader applicability in labor-intensive firms.

Fourth, we analyze the heterogeneity of structural flexibility based on the firm life cycle. Firms face varying degrees of organizational inertia across different life cycle stages. Compared to firms in mature and decline stages, those in start-up and growth stages have not yet formed entrenched vested interests, thus possessing higher organizational agility and flexibility. If pay gap convergence is indeed a rent-reduction process driven by AI technology, the estimated rent dissipation effect will inevitably be stronger in growth-stage firms—where structures are more elastic—than in mature-stage firms with rigid hierarchies. Therefore, we hypothesize that the distribution optimization effect of AI is more pronounced in early-stage firms. Following Dickinson (2011), we identify firm life cycle stages using cash flow patterns, aggregating firms in

the start-up and growth stages into the "growth group," and those in the mature, decline, and shake-out stages into the "mature group." The results in Column (4) of Table 11 indicate that the significant convergence effect of AI shocks on the intra-firm pay gap is almost entirely concentrated within the growth group, whereas it is statistically insignificant in mature-stage firms, with the inter-group difference test proving highly significant.

5.4 Excluding Alternative Explanations

The preceding mechanism tests confirm that AI shocks drive the convergence of the intra-firm pay gap primarily through the rent dissipation mechanism. However, to ensure the robustness of our findings and enhance causal inference, this section systematically rules out potential alternative explanations, specifically those driven by abnormal executive turnover, structural labor substitution, or firms' strategic catering. The corresponding empirical results are reported in Table 12.

5.4.1 Excluding Abnormal Executive Turnover and Talent Drain

The observed narrowing of the intra-firm pay gap could theoretically stem from technological shocks triggering the adverse selection or concentrated departure of highly paid managers. This would create a statistical artifact—a mechanical decline in the numerator (average executive compensation)—rather than reflecting true rent dissipation. To rule out this mechanical effect, we examine the impact of AI shocks on the voluntary executive turnover rate. Following Xu and Tan (2014) and Jiang et al. (2023), we conduct a textual analysis of the disclosed reasons for executive departures, classifying only those citing "resignation" or "personal reasons" as voluntary turnover. We scale this count by the total number of executives at the end of the previous year to construct the executive voluntary turnover rate ($ExecVoluntaryExit$).

Column (1) of Table 12 shows that the estimated coefficient of AI shocks on this turnover rate is -0.072 and is statistically insignificant. This indicates that AI shocks do not significantly disrupt top management team stability or trigger an abnormal brain

drain. Consequently, this bolsters our mechanism analysis: the reshaping of the intra-firm compensation structure is not a mere statistical byproduct of highly paid talent outflows.

5.4.2 Excluding Structural Labor Substitution and Denominator Effects

Frontier technologies like AI may directly substitute rank-and-file workers performing routine tasks. If low-paid employees are disproportionately displaced, the average compensation of the remaining rank-and-file workforce would passively rise due to the exit of low-income earners. This would statistically narrow the pay gap (i.e., a "denominator effect"). If this compositional shift dominates, the observed pay gap convergence would primarily reflect statistical changes in the labor force structure rather than a genuine top-down redistribution of positional rents.

To eliminate this alternative explanation, we construct the rank-and-file employee growth rate ($StaffGrowth$)—defined as the net increase in rank-and-file employees scaled by the total number of rank-and-file employees at the end of the previous year. As reported in Column (2) of Table 12, the estimated coefficient of AI shocks on $StaffGrowth$ is 0.595, which is statistically insignificant. This evidence suggests that, during our sample period, AI adoption did not trigger a massive contraction in the aggregate size of the rank-and-file workforce. Therefore, the asymmetric compensation dynamics cannot be explained by a denominator effect driven by structural layoffs.

5.4.3 Excluding "AI-washing" and Policy Arbitrage

The phenomenon of firms engaging in superficial or "pseudo" AI adoption poses a potential threat to our core mechanism. Firms might engage in "AI-washing"—merely inflating AI-related rhetoric in annual reports to capture external policy dividends or preferential resources—without undertaking substantive technological investments. If true, the observed compensation convergence might simply be a byproduct of other exogenous macroeconomic shocks or administrative interventions rather than genuine technological transformation.

To verify the authenticity of AI absorption, we follow Yao et al. (2024) and separately test AI word frequency ($AI_Mentions$), which proxies for strategic cognition, and AI investment level ($AIInvestment$), which proxies for substantive resource commitment. The results in Columns (3) and (4) of Table 12 demonstrate that the estimated coefficients for both variables are significantly positive. This confirms that firms not only prioritize AI applications in their strategic disclosures but also deploy substantive capital toward AI hardware and software assets. Therefore, we can confidently reject the hypothesis that our findings are driven by pseudo-AI deployment or policy arbitrage via strategic catering.

6. Conclusion and Implications

Does AI break the traditional stickiness of executive pay? We show that it does—through rent dissipation. Integrating knowledge hierarchy theory with a task automation framework, we demonstrate that AI technology reduces the cost of acquiring explicit knowledge, triggering a downward delegation of tasks. This delegation compresses executives' span of control and erodes the positional rents embedded in their compensation, while rank-and-file wages remain rigid due to implicit insurance contracts.

Empirically, utilizing a panel of China's A-share listed companies from 2015 to 2024 and exploiting the National New Generation AI Innovation and Development Pilot Zone policy as a quasi-natural experiment, we systematically test these predictions using a heterogeneity-robust imputation estimator. The results reveal that AI shocks significantly promote the convergence of the intra-firm pay gap. This convergence exhibits distinct asymmetric compensation stickiness during dynamic adjustments: executive compensation displays high downward flexibility (decreasing by approximately 6.1%), whereas rank-and-file compensation maintains stable downward rigidity. Mechanism tests indicate that this pay gap convergence fundamentally stems from the dissipation of managerial positional rents. Specifically, AI-driven capital deepening increases depreciation and amortization expenses, squeezing the

distributable surplus space via tightened financial constraints. Simultaneously, the downward delegation of tasks promotes organizational flattening, weakening management's task monopoly and bargaining power via organizational constraints. Cross-sectional heterogeneity analyses further demonstrate that this convergence effect is more pronounced in firms characterized by higher internal information asymmetry, weaker outside options for management, higher labor intensity, and those in early growth stages.

This paper makes three contributions. First, we incorporate organizational hierarchy into the task rent dissipation framework, shifting the literature's focus from horizontal task substitution to vertical structural change. Second, we provide firm-level evidence from China's A-share market on how AI shocks reshape intra-firm pay gaps—evidence that extends a literature mostly focused on macro-level skill polarization. Third, our results show that AI can serve as a governance mechanism that flattens organizations and curbs bureaucratic rent-seeking, but its high endogenous depreciation rate may trigger an "Engels' pause" pattern. This trade-off has direct implications for the design of inclusive growth policies in the AI era.

These findings should be interpreted with two caveats in mind. First, our post-shock window is short (2019–2024), and longer horizons may reveal whether the narrowing of the pay gap persists or reverses as AI capabilities evolve further. Second, our firm-level data cannot distinguish which specific workers within the hierarchy gain or lose. Access to employee-level data, patent text, or cross-country settings would sharpen our understanding of how AI reshapes distribution—and whether the "Engels' pause" pattern is a transitional phase or a permanent feature of the AI era.

REFERENCES

Acemoglu, D. and P. Restrepo, "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 2018, 108 (6), 1488-1542.

Acemoglu, D. and P. Restrepo, "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity," *The Quarterly Journal of Economics*, 2026, 141 (2), 1521-1579.

Allen, R. C., "Engels' pause: Technical change, capital accumulation, and inequality in the british industrial revolution," *Explorations in Economic History*, 2009, 46 (4), 418-435.
Azariadis, C., "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, 1975, 83 (6), 1183-1202.
Baily, M. N., "Wages and Employment under Uncertain Demand," *The Review of Economic Studies*, 1974, 41 (1), 37.
Baily, M. N., D. M. Byrne, A. T. Kane, and P. E. Soto, "Generative AI: What Type of Technology?," *AEA Papers and Proceedings*, 2026, 116, 26-30.
Bebchuk, L. A. and J. M. Fried, "Executive Compensation as an Agency Problem," *Journal of Economic Perspectives*, 2003, 17 (3), 71-92.
Bloom, D. E., K. Prettner, J. Saadaoui, and M. Veruete, "AI and the skill premium," *Finance Research Letters*, 2025, 81, 107401.
Borusyak, K., X. Jaravel, and J. Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, 91 (6), 3253-3285.
Brynjolfsson, E., D. Li, and L. Raymond, "Generative AI at Work," *The Quarterly Journal of Economics*, 2025, 140 (2), 889-942.
Chen, Dong and Ziyang Qin, "AI and Inclusive Growth: Evidence from the Use of Global Industrial Robots," *Economic Research Journal*, 2022, (4). (in Chinese)
Chen, Lin, Yuepeng Gao, and Linhui Yu, "How Does AI Change Corporate Demand for Labor? Analysis from Big Data of Recruitment Platforms," *Management World*, 2024, (6). (in Chinese)
Chu, Jian and Junxiong Fang, "Can Government Audit Restrain Executive Excess Perquisites in State-Owned Enterprises?," *Accounting Research*, 2016, (9). (in Chinese)
Core, J. E., R. W. Holthausen, and D. F. Larcker, "Corporate governance, chief executive officer compensation, and firm performance," *Journal of Financial Economics*, 1999, 51 (3), 371-406.
Dickinson, V., "Cash Flow Patterns as a Proxy for Firm Life Cycle," *The Accounting Review*, 2011, 86 (6), 1969-1994.
Eggers, A. C., G. Tuñón, and A. Dafoe, "Placebo Tests for Causal Inference," *American Journal of Political Science*, 2024, 68 (3), 1106-1121.
Garicano, L., "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, 108 (5), 874-904.
Guo, Kaiming, "AI Development, Industrial Structure Transformation and Upgrading, and Changes in Labor Income Share," *Management World*, 2019, (7). (in Chinese)
He, Xiaogang, Guoyue Zhu, and Dawei Feng, "Industrial Robot Application and Labor Income Share: Evidence from Chinese Industrial Enterprises," *China Industrial Economics*, 2023, (4). (in Chinese)
Jiang, Fuwei, Hui Ding, and Fujing Jin, "Reference Point Effect, Corporate Governance and Financial Restatement of Listed Companies," *Economic Research Journal*, 2023, (10). (in Chinese)
Jiang, Wei, Xingda Chen, and Shicheng Ni, "AI, Industrial Structure Transformation, and Crossing the 'Middle-Income Trap'," *Economic Research Journal*, 2025, (10). (in Chinese)

Kong, Dongmin, Mingli Xu, and Gaowen Kong, "Intra-Firm Pay Gap and Innovation," *Economic Research Journal*, 2017, (10). (in Chinese)
Nonaka, I., "A Dynamic Theory of Organizational Knowledge Creation," *Organization Science*, 1994, 5 (1), 14-37.
Polanyi, M., *The Tacit Dimension*, London: Routledge & Kegan Paul, 1966.
Quan, Xiaofeng, Shinong Wu, and Fang Wen, "Managerial Power, Private Benefits and Compensation Manipulation," *Economic Research Journal*, 2010, (11). (in Chinese)
Roth, J., "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, 4 (3), 305-322.
Su, Zhi, Ying Li, and Zhiyuan Cheng, "How Does AI Reshape the Factor Income Distribution Pattern? Micro Evidence Based on Firms' Labor Market Monopoly Power," *Economic Research Journal*, 2025, (9). (in Chinese)
Wang, Linhui, Shengming Hu, and Zhiqing Dong, "AI Technology, Task Attributes and Occupational Substitutability Risks: Empirical Evidence from the Micro Level," *Management World*, 2022, (7). (in Chinese)
Wang, Zhongyang and Song Tang, "The Improvement of Judicial Procedural Efficacy and Supply Chain Allocation: Evidence from the 'Separation of Complicated and Simple Cases' Reform," *The Journal of World Economy*, 2024, (9). (in Chinese)
Xu, Xixiong and Jin Tan, "Executive Compensation Contract, Reference Point Effect and Its Governance Effect: Theoretical Explanation and Empirical Evidence Based on Behavioral Economics," *Nankai Business Review*, 2014, (4). (in Chinese)
Yang, Tianyu and Yingying Zhen, "Study on the Wealth Distribution Effect of AI from a Global Perspective," *Economy and Management*, 2026, (3). (in Chinese)
Yao, Jiaquan, Kunpeng Zhang, Lipeng Guo, and Xu Feng, "How Does AI Improve Corporate Production Efficiency? From the Perspective of Labor Skill Structure Adjustment," *Management World*, 2024, (2). (in Chinese)
Yu, Li, Shuting Wang, Fan Li, and Chunrong Ai, "Labor-Saving Innovation, Monopsony and Excess Profits: Concurrently on Wage Payment Guarantee Mechanisms in the Context of Automation," *Management World*, 2025, (9). (in Chinese)
Zhang, Y. and T. Zhang, "The Economics of Digital Intelligence Capital: Endogenous Depreciation and the Structural Jevons Paradox," *arXiv Working Paper*, No. 2601.12339, 2026.
Zhao, Chunming, Yuanhao Ban, Hongbing Li, and Ye Liu, "Corporate Digital Transformation and Labor Income Share," *Journal of Finance and Economics*, 2023, (6). (in Chinese)
Zhu, Zhujun, Rancheng Xie, Hangyu Chen, Xiao Yu, and Zhifang Guo, "AI Technological Innovation and Chinese Labor Market Power: From 'Dual Growth Divergence' to Labor-Capital Compatible Development," *Management World*, 2026, (1). (in Chinese)

Table 1

Variable Definitions

Variable	Symbol	Description
A. Dependent variables		
Intra-firm pay gap	$WageGap$	The ratio of average executive compensation to average rank-and-file employee compensation.
Executive compensation	$lnAMP$	The natural logarithm of the total compensation of directors, supervisors, and senior executives divided by the number of paid directors, supervisors, and senior executives (excluding independent directors).
Rank-and-file employee compensation	$lnAEP$	The natural logarithm of total rank-and-file employee compensation divided by the number of rank-and-file employees (total employees minus the number of directors, supervisors, and senior executives). Total rank-and-file employee compensation is calculated as cash paid to and for employees plus the ending balance of payroll payable minus the beginning balance of payroll payable, less the total compensation of directors, supervisors, and senior executives.
B. Explanatory variable		
AI shocks	AI	An indicator variable that equals 1 if the city where the firm is located was approved to establish the National New Generation AI Innovation and Development Pilot Zone in the current or prior years, and 0 otherwise.
C. Control variables		
Firm size	$Size$	The natural logarithm of total assets at the end of the year.
Firm Value	$TobinQ$	Market value divided by total assets.
Profitability	ROA	Net income divided by the average balance of total assets.
Capital Structure	Lev	Total liabilities divided by total assets.
Ownership Concentration	$Top1$	The shareholding percentage of the largest shareholder.
Property Rights Nature	SOE	An indicator variable that equals 1 if the ultimate controller of the firm is a state-owned entity, and 0 otherwise.

Table 2

Descriptive Statistics

This table reports the descriptive statistics for the variables used in the analysis. Panel A presents the summary statistics for the full sample. Panel B reports the sub-sample statistics comparing the treatment and control groups, along with the test of mean differences. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. The mean difference test employs an independent two-sample t-test with heteroskedasticity-robust standard errors.

Notes: While Panel A reports summary statistics for the full unbalanced panel (2015–2024), Panel B reports the mean difference tests utilizing the pre-treatment baseline averages (2015–2017) to evaluate initial structural endowments prior to the policy shock.

Panel A：Full Sample

Variable	Observations	Mean	S.D.	Min	Max
$WageGap$	28723	4.842	3.461	1.207	21.574
$lnAMP$	28723	13.193	0.618	11.810	14.954
$lnAEP$	28723	11.798	0.461	10.783	13.011
AI	28723	0.245	0.430	0.000	1.000
$Size$	28723	22.122	1.265	19.969	26.012
$TobinQ$	28723	2.408	1.360	0.982	8.993
ROA	28723	0.043	0.046	-0.109	0.179
Lev	28723	0.402	0.194	0.066	0.863
$OwnCon$	28723	34.798	14.199	9.686	73.276
SOE	28723	0.328	0.467	0.000	1.000

Panel B: Pre-treatment Characteristics and Mean Difference Tests

	Treatment Group		Control Group		
Variable	Mean	S.D.	Mean	S.D.	Mean Difference
$Size$	22.202	(1.342)	22.009	(1.186)	0.193***
$TobinQ$	2.508	(1.469)	2.377	(1.378)	0.131**
ROA	0.043	(0.045)	0.041	(0.049)	0.003
Lev	0.416	(0.197)	0.396	(0.195)	0.020***
$OwnCon$	35.264	(14.679)	34.059	(13.777)	1.205**
SOE	0.351	(0.476)	0.285	(0.449)	0.066***
Number of firms	1418		1743		

Table 3

Baseline Regression Results: AI Shocks and the Reshaping of Intra-firm Compensation Structure

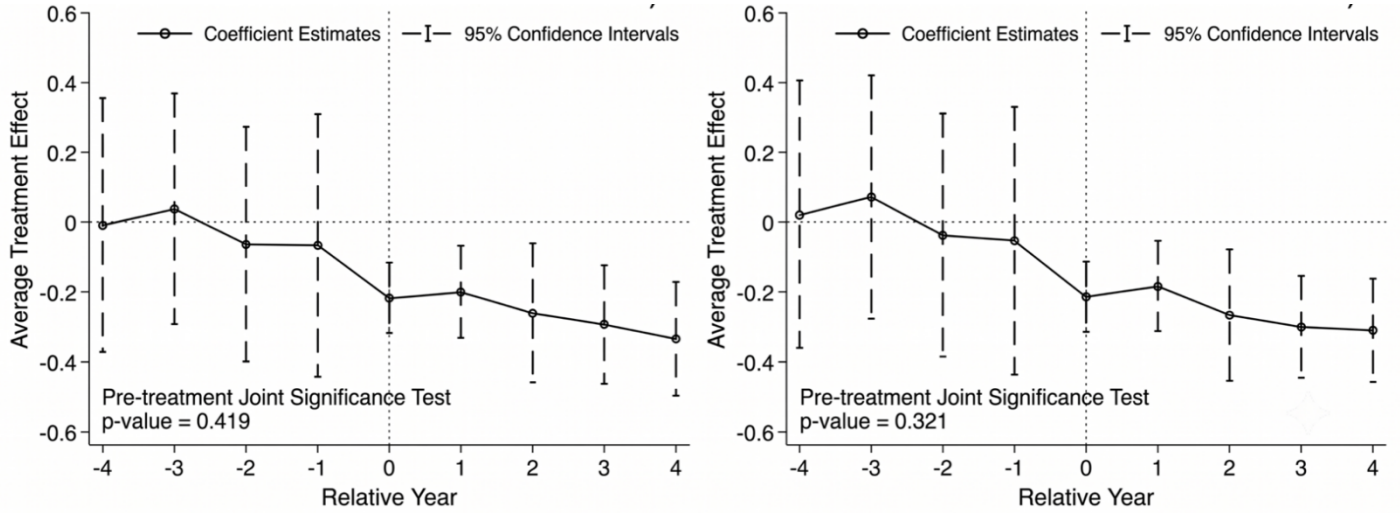
This table presents the baseline regression results for the impact of AI shocks on the internal compensation structure. Columns (1) and (2) report the effects on the intra-firm pay gap ($WageGap$), while Columns (3) and (4) report the effects on the executive compensation level ($lnAMP$) and rank-and-file compensation level ($lnAEP$), respectively. All estimations are conducted using the robust imputation estimator proposed by Borusyak et al. (2024). "Pre-treatment characteristics time trends" refer to the interactions between linear time trends and the pre-policy baseline averages (2015–2017) of the firm-level control variables. Robust standard errors clustered at the prefecture-level city are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Variable	（1）	（2）	（3）	（4）
	$WageGap$	$WageGap$	$lnAMP$	$lnAEP$
AI	-0.318***	-0.291***	-0.061***	-0.004
	(0.065)	(0.053)	(0.008)	(0.008)
Pre-treatment characteristics time trends	No	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	28723	28723	28723	28723

Figure 1

Dynamic Treatment Effects and Parallel Trends Test (Imputation Estimator)

This figure plots the dynamic average treatment effects of AI shocks on the intra-firm pay gap ($WageGap$) based on the imputation estimator (Borusyak et al., 2024). The vertical dashed line at $k = 0$ marks the period immediately preceding the official implementation of the policy (capturing the anticipation effect via the shift(1) specification). The horizontal dashed line denotes the zero-effect baseline. The error bars represent 95% confidence intervals derived from robust standard errors clustered at the prefecture-level city. The p-values located at the bottom left of each panel correspond to the Wald joint significance test for the strictly pre-treatment periods ($k \in [-4, -1]$).



Left Panel: Intra-firm Pay Gap (Excluding Linear Time Trends of Pre-treatment Characteristics)

Right Panel: Intra-firm Pay Gap (Including Linear Time Trends of Pre-treatment Characteristics)

Table 4

Placebo Outcome Test

This table reports the results of the placebo outcome tests following the framework of Eggers et al. (2024). We replace the true dependent variable with pseudo-outcomes that theoretically satisfy the No Average Treatment Effect (NATE) assumption: Average Executive Age, Number of Shareholders, Effective Tax Rate, and Ratio of Non-operating Income. All specifications are estimated using the robust imputation estimator (Borusyak et al., 2024). "Pre-treatment characteristics time trends" refer to the interactions between linear time trends and the pre-policy baseline averages (2015–2017) of the firm-level control variables. Robust standard errors clustered at the prefecture-level city are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Variable	（1）	（2）	（3）	（4）
	Executive Average Age	Number of Shareholders	Effective Tax Rate	Non-operating Income Ratio
AI	0.060	0.001	0.005	-0.004
	(0.116)	(0.024)	(0.005)	(0.007)
Pre-treatment characteristics time trends	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	28633	28721	28488	28546

Figure 2

Kernel Density Distribution of Placebo Treatment Effects

This figure plots the kernel density distribution of the pseudo-treatment effect coefficients generated from 500 Monte Carlo simulation regressions. In each iteration, the true policy implementation years are randomly reshuffled across cities to construct spatiotemporally misaligned pseudo-treatments, and the model is re-estimated using the imputation estimator (Borusyak et al., 2024). The solid curve represents the empirical distribution of the simulated coefficients, which is strictly centered around zero. The vertical dashed line denotes the true baseline regression coefficient (-0.291) from Table 3, Column (2), which falls squarely in the extreme left tail of the placebo distribution, thereby demonstrating the robustness of the causal identification.

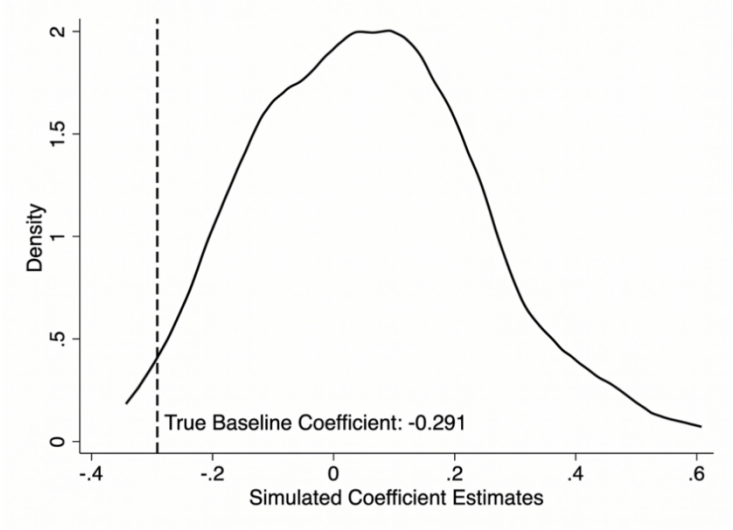


Table 5

Placebo Population Test

This table reports the results of the placebo population tests (Eggers et al., 2024). To verify that the baseline results are not driven by unobservable macroeconomic common trends, we re-estimate the model using three theoretically immune subpopulations: a State-Owned Enterprise (SOE) subsample, a technologically immune traditional-industry subsample, and a low-technology dependence subsample. The dependent variable across all columns is the intra-firm pay gap ($WageGap$). All specifications are estimated using the robust imputation estimator (Borusyak et al., 2024). "Pre-treatment characteristics time trends" refer to the interactions between linear time trends and the pre-policy baseline averages (2015–2017) of the firm-level control variables. Robust standard errors clustered at the prefecture-level city are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Variable	（1）	（2）	（3）
	$WageGap$		
	SOE Sample	Technology-Immune Sample	Low-Tech-Dependent Sample
AI	-0.138	0.017	-0.092
	(0.114)	(0.184)	(0.150)
Pre-treatment characteristics time trends	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	9723	2271	7745

Table 6

Robustness Checks: Excluding the Interference of Contemporaneous Policies

This table reports robustness checks addressing potential confounding from contemporaneous policies. Column (1) excludes state-owned enterprises (SOEs) to prevent interference from the SOE executive compensation cap mandate. Column (2) utilizes the full sample but incorporates ownership-by-year joint fixed effects to absorb macroeconomic shocks specific to property rights. Columns (3) and (4) isolate the baseline treatment effect by explicitly treating the National AI Innovative Application Leading Zones and National Big Data Comprehensive Pilot Zones as time-varying confounders in the first stage of the imputation, respectively. The dependent variable is the intra-firm pay gap ($WageGap$). All specifications are estimated using the robust imputation estimator (Borusyak et al., 2024).

Variable	（1）	（2）	（3）	（4）
	$WageGap$			
	Non-SOE Sample	Full Sample	AI Leading Zones	Big Data Pilot Zones
AI	-0.395***	-0.296***	-0.156***	-0.312***
	(0.089)	(0.056)	(0.051)	(0.054)
AI Pilot Zones	No	No	Yes	No
Big Data Pilot Zones	No	No	No	Yes
Pre-treatment characteristics time trends	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Ownership-Year FE	No	Yes	No	No
Observations	18161	28257	28723	28723

Table 7

Additional Robustness Checks

This table presents further robustness checks for the baseline results. Column (1) mitigates firm-specific demand shock interference by restricting the sample to firms exhibiting ex-ante demand expansion. Column (2) controls for time-varying regional unobservables by interacting baseline city-level macro characteristics with linear time trends. Column (3) eliminates nominal inflation effects by deflating all financial variables using a 2015 fixed-base Consumer Price Index (CPI). Column (4) mitigates spatial self-selection bias by excluding firms that changed their registered prefecture-level city during the sample period. Column (5) addresses compositional bias by restricting the analysis to a strongly balanced panel. The dependent variable across all columns is the intra-firm pay gap ($WageGap$). All specifications are estimated using the robust imputation estimator (Borusyak et al., 2024).

Variable	（1）	（2）	（3）	（4）	（5）
	$WageGap$				
	Demand Expansion Sample	Regional Unobservables	CPI Deflated	Excluding Relocation	Strongly Balanced Panel
AI	-0.318***	-0.323***	-0.290***	-0.337***	-0.277***
	(0.060)	(0.099)	(0.053)	(0.064)	(0.071)
Baseline pre-treatment characteristics time trends	Yes	Yes	Yes	Yes	Yes
City pre-treatment characteristics time trends	No	Yes	No	No	No
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	22377	28420	28723	27718	19640

Table 8

Excluding the Interference of Major Public Health Emergency

This table reports robustness checks addressing the potential confounding effects of the global COVID-19 pandemic. Column (1) strictly excludes firms located in high-risk cities. Column (2) controls for heterogeneous pandemic exposure by interacting the pre-pandemic geographic distance to Wuhan and the city's population density with linear time trends. Column (3) conducts a time-truncation test by splitting the treatment effect into pre-pandemic (τ_{pre}, before 2020) and post-pandemic (τ_{post}, 2020 and beyond) phases to verify temporal continuity. The dependent variable is the intra-firm pay gap ($WageGap$). All specifications are estimated using the robust imputation estimator (Borusyak et al., 2024).

Variable	（1）	（2）	（3）
	$WageGap$		
	Excluding High-Risk Cities	Controlling for Pandemic Exposure	Time-Truncation Test
AI	-0.241***(0.062)	-0.254***(0.082)	
$AI(\tau_{post})$ (2020 and after)			-0.321***(0.065)
$AI(\tau_{pre})$ (Before 2020)			-0.196***(0.060)
Difference ($\tau_{post} - \tau_{pre}$)			-0.124(0.084)
Baseline pre-treatment characteristics time trends	Yes	Yes	Yes
Pandemic exposure time trends	No	Yes	No
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	17384	28427	28723

Table 9

Test Results of the Rent Dissipation Mechanism

This table presents the tests for the rent dissipation mechanism. Columns (1) and (2) examine the effect of AI on executive excess compensation ($ExcessPay$) and excess perquisites ($ExcessPerks$), respectively. Columns (3) and (4) test the impact of $ExcessPay$ and $ExcessPerks$ on the intra-firm pay gap ($WageGap$). Executive excess compensation in Column (1) is the residual extracted from the following model: $LnAMP_{i,t} = \alpha_0 + \alpha_1 LnSize_{i,t} + \alpha_2 ROA_{i,t} + \alpha_3 TobinQ_{i,t} + \alpha_4 Lev_{i,t} + \sum Industry + \sum Year + \varepsilon_{i,t}$. Executive excess perquisites in Column (2) is the residual extracted from the following model: $\frac{Perks_{i,t}}{Asset_{i,t-1}} = \alpha_0 + \alpha_1 \frac{1}{Asset_{i,t-1}} + \alpha_2 \frac{\Delta Sales_{i,t}}{Asset_{i,t-1}} + \alpha_3 \frac{PPE_{i,t}}{Asset_{i,t-1}} + \alpha_4 \frac{Inventory_{i,t}}{Asset_{i,t-1}} + \alpha_5 \ln(Emp_{i,t}) + \sum Industry + \sum Year + \varepsilon_{i,t}$. $Perks$ represents the absolute amount of implicit perquisites, calculated as administrative expenses minus the total compensation of directors, supervisors, and senior management, and the amortization of intangible assets.

Variable	（1）	（2）	（3）	（4）
	$ExcessPay$	$ExcessPerks$	$WageGap$	$WageGap$
AI	-0.046***(0.008)	-0.008***(0.002)		
$ExcessPay$			4.539***(0.079)	
$ExcessPerks$				0.519**(0.211)
Pre-treatment characteristics time trends	Yes	Yes	No	No
Time-varying control variables	No	No	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	28268	27636	27884	26931

Table 10

Suggestive Evidence on the Rent Dissipation Mechanism

This table provides further suggestive evidence for the rent dissipation mechanism. The dependent variables are the log of depreciation and amortization ($DepAmort$), the ratio of deputy executives ($DeputyRatio$), the office administrative expense ratio ($OfficeExpRatio$), and the travel administrative expense ratio ($TravelExpRatio$). Depreciation and amortization in Column (1) is calculated as the natural logarithm of the sum of fixed asset depreciation and intangible asset amortization.

Variable	（1）	（2）	（3）	（4）
	$DepAmort$	$DeputyRatio$	$OfficeExpRatio$	$TravelExpRatio$
AI	0.049*	-0.824**	-0.033***	-0.021***
	(0.026)	(0.398)	(0.010)	(0.006)
Pre-treatment characteristics time trends	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	28540	28723	24644	20898

Table 11

Heterogeneity Analysis

This table presents the heterogeneity analysis regarding the effect of AI shocks on the intra-firm pay gap ($WageGap$) across different firm characteristics. Panel A reports the average treatment effects estimated via the imputation estimator for the high (τ_{High}) and low (τ_{Low}) groups partitioned by four dimensions: Internal Information Environment, Executive Outside Options, Capital Intensity, and Firm Life Cycle. Panel B formally tests the differences between these groups using a triple-differences framework. To rigorously ensure parallel trends across sub-dimensions, the models control for group-specific linear time trends, constructed as the interaction between the respective grouping dummy and a continuous year variable.

Variable	（1）	（2）	（3）	（4）
	$WageGap$			
	Internal Information Environment	Executive Outside Options	Capital Intensity	Firm Life Cycle
Panel A：Average Treatment Effects by Heterogeneous Groups				
High Group Treatment Effect (τ_{High})	-0.565***(0.118)	-0.192***(0.060)	-0.190**(0.079)	-0.098(0.101)
Low Group Treatment Effect (τ_{Low})	-0.239***(0.053)	-0.403***(0.064)	-0.324***(0.054)	-0.399***(0.063)
Panel B：Triple-differences Test				
Group Difference ($\tau_{High} - \tau_{Low}$)	-0.325***(0.114)	0.211***(0.066)	0.134*(0.073)	0.301**(0.119)
Pre-treatment characteristics time trends	Yes	Yes	Yes	Yes
Group × Year trends	Yes	Yes	Yes	Yes
Firm / Year FE	Yes	Yes	Yes	Yes
Observations	28481	28713	28713	28723

Table 12

Tests Excluding Alternative Competing Explanations

This table reports the empirical results of tests designed to exclude potential alternative explanations for the convergence of the intra-firm pay gap. Column (1) examines the impact of AI shocks on the executive voluntary turnover rate ($ExecVoluntaryExit$) to rule out abnormal talent drain. Column (2) tests the rank-and-file employee growth rate ($StaffGrowth$) to exclude denominator effects driven by structural labor substitution. Columns (3) and (4) test AI word frequency ($AI_Mentions$) and AI investment level ($AIInvestment$), representing strategic cognition and substantive resource input, respectively, to rule out strategic catering and pseudo-AI deployment.

Variable	（1）	（2）	（3）	（4）
	$ExecVoluntaryExit$	$StaffGrowth$	$AI_Mentions$	$AIInvestment$
AI	-0.072	0.595	0.090***	0.047**
	(0.134)	(0.531)	(0.023)	(0.023)
Pre-treatment characteristics time trends	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	28031	27737	28720	28719

Appendix 1: Dynamic Effect Analysis and Identification of Policy Anticipation Effects

When evaluating the impact of the National New Generation AI Innovation and Development Pilot Zone policy on the intra-firm pay gap, verifying the parallel trends assumption is the core prerequisite for valid causal identification in a difference-in-differences (DiD) framework. To rigorously examine this assumption and the dynamic evolutionary process of the policy, we first plot the dynamic effects using a naive event-study specification that fails to account for the policy anticipation effect (i.e., mechanically defining the official policy approval year as the treatment onset, $t = 0$) (Appendix Figure 1). The corresponding static baseline regression results for this unshifted specification are reported in Appendix Table 1.

Observing the graphical trajectory intuitively reveals that under the naive specification (without intervention on the policy timing), the estimated coefficients for the intra-firm pay gap in the period immediately preceding the policy (relative year $t - 1$) exhibit a stark downward deviation. A rigorous statistical test further corroborates this observation: the pre-treatment joint significance test strongly rejects the null hypothesis that the pre-treatment dynamic coefficients are jointly zero ($p = 0.000$). However, an examination of the dynamic trajectory in Appendix Figure 1 confirms that this significant pre-treatment deviation does not originate from a long-term structural trend difference—the coefficients for $t - 2$ and earlier periods remain stable and bound tightly around zero. Instead, the deviation is highly concentrated exclusively in the period immediately preceding the policy ($t - 1$). This characteristic indicates that, given the extensive administrative incubation period required for the application and approval of national pilot zones, treated firms formulated policy expectations prior to the official rollout and subsequently made forward-looking adjustments to their intra-firm compensation structures. It is precisely this substantive, micro-level anticipatory response concentrated in period $t - 1$ that drives the failure of the pre-treatment joint significance test.

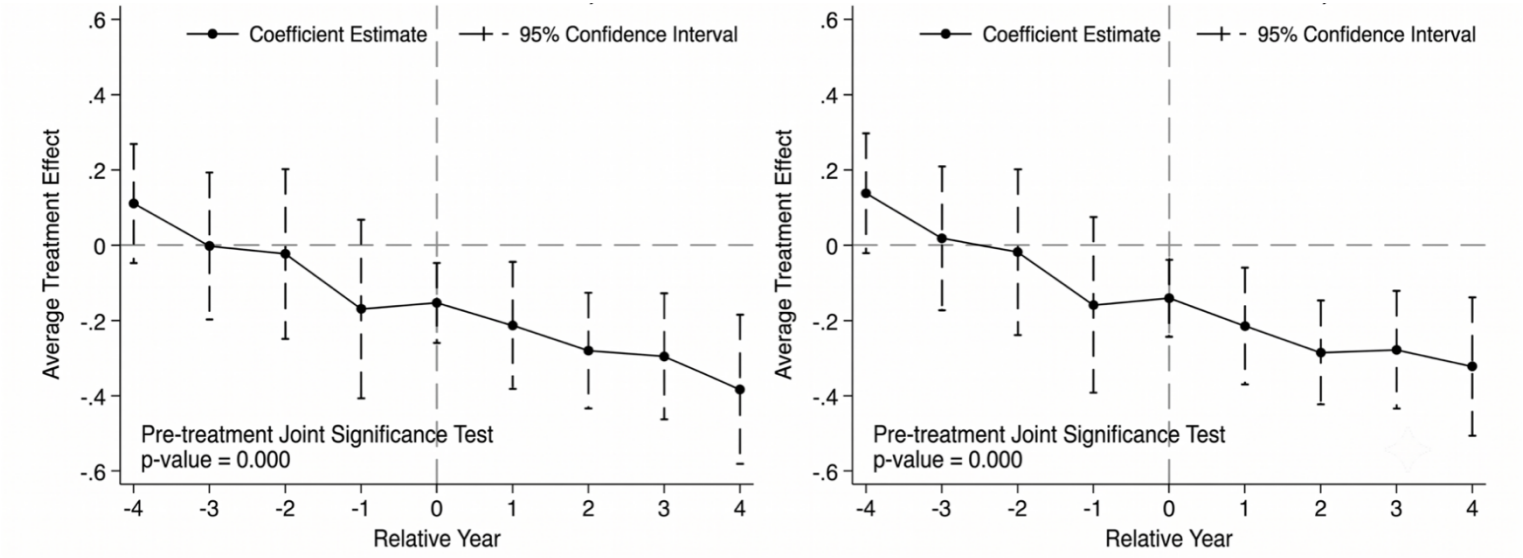
In the classic panel event-study framework, Freyaldenhoven et al. (2019) emphasize that the manifestation of pre-trends does not invariably imply a fundamental violation of the parallel trends assumption; rather, it frequently stems from the endogenous anticipation of policy shocks by market participants. As Malani and Reif (2015) point out, such forward-looking expectations often trigger significant behavioral responses before a policy is officially enacted. In this context, if the year of official policy approval is mechanically defined as the treatment time, the substantive anticipatory responses of micro-firms occurring in period $t-1$ are erroneously categorized into the untreated baseline period, resulting in severe control group contamination. Although Appendix Table 1 demonstrates that AI shocks still exert a significantly negative effect on the intra-firm pay gap (Columns 1 and 2) and executive compensation (Column 3) without accounting for anticipation, the erroneous inclusion of the contaminated $t-1$ sample—which already embeds wage-compression expectations—into the pre-treatment reference frame artificially depresses the counterfactual baseline. This inevitably causes the model to suffer from attenuation bias, distorting the true dynamic trajectory and net causal effect of the policy.

To correct this identification bias, our baseline model in the main text applies a one-period lead to the policy treatment timing under the heterogeneity-robust estimation framework of Borusyak et al. (2024)—namely, introducing the shift(1) specification. Grounded in the logic of economic behavior, this strategy systematically incorporates the anticipatory responses of micro-level firms in the period preceding the shock into the treatment effect evaluation window. By reconstructing a pristine counterfactual baseline, it effectively absorbs pre-treatment dynamic interference. The parallel trends test in the main text (Figure 1) confirms that after rigorously identifying and absorbing the policy anticipation effect in period $t-1$, the pre-treatment joint significance test is successfully passed, strictly satisfying the parallel trends assumption.

Appendix Figure 1

Dynamic Treatment Effects and Parallel Trends Test (Without Accounting for Policy Anticipation)

This figure plots the dynamic average treatment effects on the intra-firm pay gap ($WageGap$) under a naive specification that mechanically defines the official policy approval year as the treatment onset ($t = 0$), thereby failing to account for policy anticipation. The p-values at the bottom left of each panel correspond to the Wald joint significance test for the pre-treatment periods ($k \in [-4, -1]$), demonstrating a stark rejection of the parallel trends assumption due to the anticipatory contamination of period $t - 1$.



Left Panel: Intra-firm Pay Gap (Excluding Linear Time Trends of Pre-treatment Characteristics)

Right Panel: Intra-firm Pay Gap (Including Linear Time Trends of Pre-treatment Characteristics)

Appendix Table 1

Baseline Regression Results (Without Accounting for Policy Anticipation Effects)

This table presents the baseline regression results estimated without accounting for the policy anticipation effect. Columns (1) and (2) report the effects on the intra-firm pay gap ($WageGap$), while Columns (3) and (4) report the effects on the executive compensation level ($lnAMP$) and rank-and-file compensation level ($lnAEP$), respectively. All estimations employ the imputation estimator (Borusyak et al., 2024). The regression sample size (30,380) under this naive specification (i.e., without introducing the shift(1) parameter) is larger than that of the baseline model (28,723). This discrepancy primarily arises from relaxed temporal constraints during the identification window: under the unshifted specification, the extraction window for the predetermined control variables is extended to 2018, and treated firms benefit from a longer effective pre-treatment observation period. This retains a subset of short-panel firms that entered the sample later or possessed shorter pre-treatment series. Robust standard errors clustered at the prefecture-level city are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Variable	（1）	（2）	（3）	（4）
	$WageGap$	$WageGap$	$lnAMP$	$lnAEP$
AI	-0.292***	-0.271***	-0.055***	-0.005
	(0.072)	(0.061)	(0.010)	(0.006)
Pre-treatment characteristics time trends	No	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	30380	30380	30380	30380

Appendix 2: Unbalanced Panel Data Structure and Observation Period Distribution

To further validate the robustness of our data screening procedures and the suitability of the resulting dataset for the imputation estimator, Appendix Table 2 reports the structural characteristics of the unbalanced panel utilized in our final regression sample. As detailed in Panels A and B, the refined sample comprises 3,161 unique listed firms. The minimum valid observation span for any single firm is exactly two years, while both the median and the 75th percentiles reach the full ten-year sample horizon. This confirms that we have successfully eliminated singleton observations that preclude the identification of unit fixed effects and aggregate period fixed effects, thereby ensuring that every retained firm-level time series exhibits sufficient longitudinal variation.

Turning to the distribution of panel observation patterns in Panel C, firms with a strictly balanced panel—defined by complete observations across all 10 consecutive years—constitute 62.13% (1,964 firms) of the sample. This substantial core provides ample degrees of freedom to reliably estimate aggregate macroeconomic time trends. Furthermore, the unbalanced panel explicitly incorporates late-listing firms characterized by deferred entry (e.g., the "0111111111" and "0011111111" patterns), alongside firms exhibiting natural attrition or delisting (e.g., the "1111111110" pattern). Preserving this natural corporate life cycle not only maximizes the retention of informational content—a distinct methodological advantage of the state-of-the-art imputation estimator—but also effectively mitigates the survivorship bias that is inherently introduced when artificially truncating data to enforce a strongly balanced panel.

Appendix Table 2

Structural Characteristics and Distribution of the Unbalanced Panel Data

This table details the structural characteristics of the unbalanced panel dataset utilized in our primary empirical analyses. Panel A outlines the fundamental dimensions of the panel. Panel B reports the distributional percentiles of the observation period length (in years) per unique firm. Panel C enumerates the top ten most frequent observation patterns across the sample. Note: Within the observation patterns denoted in Panel C, a value of "1" designates the presence of a valid firm-year observation, whereas a "0" designates missing data for that specific year. The comprehensive sample period spans 10 consecutive effective periods (2015–2024).

Panel A：Fundamental Panel Dimensions

Indicator	Value	Description
Sample window	2015-2024	Total of 10 effective periods
Unique firm count	3,161	Total number of distinct listed firms

Panel B：Distribution of Observation Period Length (Years per Firm)

Min	5th	25th	Median	75th	95th	Max
2	5	9	10	10	10	10

Panel C：Top 10 Dominant Panel Observation Patterns

No.	Frequency	Percentage	Cumulative Percentage	Observation Pattern
1	1964	62.13%	62.13%	1111111111
2	288	9.11%	71.24%	0111111111
3	230	7.28%	78.52%	0011111111
4	70	2.21%	80.73%	1111111110
5	35	1.11%	81.84%	1111000000
6	33	1.04%	82.89%	1111101111
7	30	0.95%	83.83%	1011111111
8	28	0.89%	84.72%	1111100000
9	28	0.89%	85.61%	1111111100
10	27	0.85%	86.46%	1111111000
Others	428	13.54%	100.00%	(Other structural patterns)
Total:	3161	100.00%	100.00%	