

Artificial Intelligence Adoption and the Within-Firm Labor Share: Evidence from Chinese Listed Firms

Jiawei Ma

p2626244@mpu.edu.mo

MPI GTR Center: Macao Polytechnic University Gaming Teaching and Research Center

<https://orcid.org/0009-0000-6423-3039>

Research Article

Keywords: artificial intelligence, labor share, automation, markups, monopsony, China

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11368916/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Artificial Intelligence Adoption and the Within-Firm Labor Share: Evidence from Chinese Listed Firms

Jiawei Ma, Ph.D.

Faculty of Management Sciences, Macao Polytechnic University

Macao SAR, China

E-mail: P2626244@mpu.edu.mo | ORCID: 0009-0000-6423-3039

Abstract

Does artificial intelligence rewire the division of value added between labor and capital inside the firm? We answer this question using a staggered-adoption panel of 2,274 Chinese A-share listed firms over 2010–2024 (26,579 firm-year observations), measuring AI adoption with a Chinese-BERT classifier applied to annual-report text. Identification combines a two-way fixed-effects design, heterogeneity-robust estimators, and an instrument based on provincial program-controlled telephone switch capacity density in 1992. AI adoption lowers the within-firm labor share by 1.18 percentage points (2.8 percent of the sample mean). The decline reflects decoupling rather than immiseration: value added per worker rises 5.8 percent and average wages rise 2.1 percent while employment falls 3.2 percent—wages fail to keep pace with productivity. The wage-bill share of operating costs falls 0.92 percentage points and markups rise by 0.035. Estimates are stable across Callaway–Sant’Anna, Sun–Abraham, and Borusyak–Jaravel–Spiess estimators, a sixteen-specification robustness battery, and IV-2SLS (10 percent level with provincial controls). Task substitution, productivity-and-scale, and markup channels descriptively account for 97.5 percent of the effect, which is concentrated in labor-intensive firms.

Keywords: artificial intelligence; labor share; automation; markups; monopsony; China

JEL Classification: J23, J24, J31, L25, O33, E25

1. Introduction

For much of the twentieth century, the division of national income between labor and capital was stable enough to count as one of Kaldor’s stylized facts. That stability has ended. The corporate labor share has fallen by roughly five percentage points across countries since the mid-1970s, and within-country declines dominate the global pattern (Karabarbounis and Neiman 2014). In the United States, the decline is concentrated in the payroll component of value added (Elsby, Hobijn, and Şahin 2013) and coincides with a rising pure-profit share, especially in large firms (Barkai 2020). China displays the same phenomenon through a different institutional lens: official statistics record a fall and partial recovery of the aggregate labor share whose magnitude depends sensitively on measurement conventions (Bai and Qian 2009; Luo and Zhang 2009; Liu, Shen, and Yao 2022). Explaining this decline—and identifying which forces operate inside firms rather than between them—is one of the central open questions in labor and macroeconomics (Karabarbounis 2024).

Artificial intelligence is a natural suspect. In the task-based framework, automation shifts the frontier of tasks performed by capital into tasks previously performed by labor, and the equilibrium labor share depends on the race between this displacement effect and countervailing productivity and reinstatement effects (Acemoglu and Restrepo 2018, 2019). Unlike industrial robots, which substitute for manual tasks, contemporary AI performs prediction, language, and coordination tasks, and its exposure reaches high up the wage distribution (Eloundou et al. 2024). Whether an adopting firm then pays its remaining workers a constant share of a larger pie, or a smaller share, determines whether AI arrives as a complement to labor’s income or as a quiet redistribution toward capital and profits. This paper asks that question at the level where value added is actually divided: when a firm adopts AI, what happens to its own labor share?

Three gaps in the literature motivate the question. First, the leading firm-level explanation of the labor-share decline is reallocation, not within-firm change. Autor et al. (2020) show that the fall of the U.S. labor share is dominated by the market-share expansion of low-labor-share “superstar” firms, and Kehrig and Vincent (2021) document the same compositional anatomy in U.S. manufacturing; in both accounts the typical firm’s own labor share moves little or even rises. What these designs cannot show is whether a discrete technology adoption changes the labor share inside the adopting firm—the within-firm margin is differenced out. We estimate exactly that margin. Second, the automation evidence is industry- and region-level and divided in ways that bear directly on the sign of the within-firm effect: robots reduce local employment and wages in the United States (Acemoglu and Restrepo 2020) but leave German incumbents’ wages largely untouched (Dauth et al. 2021), while raising productivity without reducing aggregate employment in cross-country industry panels (Graetz and Michaels 2018). Firm-level AI investments appear to drive growth through product innovation rather than labor substitution (Babina et al. 2024), and vacancy-based measures show large skill-composition shifts but small net employment effects

(Acemoglu et al. 2022). None of these studies measures the division of value added. Third, little is known about settings with high labor intensity, segmented internal labor markets, and heterogeneous employer market power—conditions under which wage setting, and hence the labor share, responds to technology in theoretically distinct ways (Manning 2003, 2021; Card et al. 2018). China offers this setting, and its evidence on robots and factor shares shows that automation can compress the within-firm labor share through capital-labor substitution and wage-rate channels (He, Zhu, and Feng 2023), most starkly in emerging superstar firms (Chen, Yao, and Zheng 2024; Wang and Dong 2020).

China is also a natural laboratory for measurement. A-share listed firms file standardized annual reports whose management-discussion sections describe technology deployment in unusual detail, and recent work validates text-based firm-year measures of AI adoption on exactly this corpus (Wang and Qiu 2025; Chen, Yang, and Wu 2025), in the tradition of firm-level text measurement (Hassan et al. 2019; Gentzkow, Kelly, and Taddy 2019). We fine-tune a Chinese-BERT classifier to detect substantive AI deployment. It attains an out-of-sample F1 of 0.91 and Cohen’s κ of 0.81 against hand labels, and its firm-year scores correlate at 0.72–0.79 (Spearman) with three external benchmarks. Applied to reports from 2010 to 2024, it yields a staggered adoption panel of 1,286 adopting firms in cohorts from 2013 to 2021—with the mode cohort of 218 firms in 2018—and 988 never-treated firms (43.4 percent of the sample of 2,274). For causal identification we exploit the geography of China’s 1992 program-controlled telephone switching network: provincial switch capacity density, fixed in 1992 and interacted with the post-2013 diffusion period, predicts adoption with a first-stage coefficient of 0.1876 and an F-statistic of 690.6, while predating all sample outcomes by nearly two decades.

The empirical design has three layers. The baseline is a two-way fixed-effects specification with firm and industry-by-year fixed effects. Because adoption is staggered, we validate the baseline with the Callaway–Sant’Anna, Sun–Abraham, and Borusyak–Jaravel–Spiess estimators, which all lie within ± 10 percent of the baseline, and with a joint pre-trend test that does not reject ($p = 0.4335$). We then instrument adoption with the 1992 switch capacity density; the 2SLS estimate is 1.22 times the baseline (-1.4394 versus -1.1800) and remains stable in magnitude, at 1.10 times the baseline, once 1992 provincial controls are added (significant at the 10 percent level). A sixteen-specification robustness battery, including a placebo adoption event among never-treated firms (coefficient $+0.3595$, $t = 1.10$), supports the design.

We document six facts. AI adoption reduces the within-firm labor share by 1.1800 percentage points ($SE = 0.1475$, $t = -8.00$), about 2.8 percent of the sample mean of 42.60. The decomposition of this decline is the paper’s central finding. Employment falls by 0.0320 log points ($SE = 0.0050$, $t = -6.40$), yet the average wage rises by 0.0210 log points ($SE = 0.0043$, $t = 4.91$) and value added per worker rises by 0.0580 log points ($SE = 0.0081$, $t = 7.20$). The labor share falls, in other words, not because workers are impoverished but because wages decouple from productivity: a 5.8 percent productivity gain is met by a 2.1

percent wage gain. Consistent with a distributional shift, the wage-bill share of operating costs falls by 0.9200 percentage points ($SE = 0.1769$, $t = -5.20$; 3.2 percent of the mean of 28.50) and the estimated price-cost markup rises by 0.0350 ($SE = 0.0051$, $t = 6.80$; 3.0 percent of the mean of 1.18). A formal decomposition attributes 30.0 percent of the labor-share decline to task substitution, 38.5 percent to the productivity-and-scale channel, and 29.0 percent to the markup channel—descriptively 97.5 percent jointly. The effect concentrates where the theory predicts: in labor-intensive firms, the top-to-bottom quartile contrast is 2.60 to 1, larger than the contrasts by ownership (1.90), sector (1.80), technology intensity (1.60), or size (2.20), all significant at the 5 percent level.

The paper makes three contributions. First, it provides the first firm-level causal evidence on how AI adoption changes the within-firm labor share, complementing the reallocation account of Autor et al. (2020) and Kehrig and Vincent (2021) with the margin they difference out. Our findings show the two accounts are consistent: adoption lowers the adopter’s own labor share even as aggregate declines are amplified by the expansion of such firms. Second, it adjudicates the displacement-versus-productivity debate for AI specifically (Acemoglu and Restrepo 2018; Graetz and Michaels 2018; Babina et al. 2024): the productivity effect dominates employment and wage margins—our adopters are not running “so-so automation” (Acemoglu and Restrepo 2019)—yet the labor share still falls because productivity gains are not passed through to wages, a signature of imperfect labor- and product-market competition (Manning 2003; Card et al. 2018; De Loecker, Eeckhout, and Unger 2020). Third, methodologically, it pairs a validated text-based adoption measure with the modern staggered-DID toolkit and a historically motivated instrument in the world’s largest emerging stock market, extending the A-share text-measurement paradigm (Wang and Qiu 2025; Chen, Yang, and Wu 2025) from employment and cost outcomes to the functional distribution of income, and answering the call for micro evidence on single drivers of the labor share (Karabarbounis 2024).

The remainder of the paper proceeds as follows. Section 2 develops the theoretical framework and Section 3 states four hypotheses. Section 4 describes the data and variable construction. Section 5 sets out the empirical strategy. Section 6 reports baseline and robust estimates, Section 7 the causal-identification battery, Section 8 the mechanism decomposition, and Section 9 heterogeneity. Section 10 presents a causal machine-learning extension, and Section 11 concludes.

2. Theoretical Framework

This section organizes the forces through which AI adoption can move the within-firm labor share. We first describe the task-based shock and its two opposing effects on factor shares, then the product-market markup channel, then the labor-market wage-setting channel, and finally a compact model that delivers four testable predictions.

2.1 Task-Based Automation: Displacement versus Productivity Effects

We view AI as capital that performs tasks. Automation expands the set of tasks in which capital can substitute for labor; this displacement effect reduces labor demand and, holding the capital share's own determinants fixed, reduces the labor share (Acemoglu and Restrepo 2018). The same framework supplies the counterweight: automation raises total factor productivity, and the induced scale expansion raises demand for labor in the tasks that remain, while new-task creation reinstates labor directly (Acemoglu and Restrepo 2018, 2019). The sign of the labor-share response is therefore not a technological constant. It depends on the elasticity of substitution between AI capital and labor and on how completely productivity gains are transmitted to wages. A useful benchmark is what Acemoglu and Restrepo (2019) call “so-so automation”: technologies that displace labor while delivering only trivial productivity gains unambiguously lower the labor share. Genuine productivity-enhancing AI is ambiguous in principle—the labor share can survive a displacement effect if wages rise one-for-one with output per worker—and the empirical question is the degree of pass-through.

The existing evidence frames the margins. Robots raise labor productivity and TFP without reducing aggregate employment in industry panels (Graetz and Michaels 2018), but reduce employment and wages in exposed local labor markets (Acemoglu and Restrepo 2020) and task-displaced demographic groups (Acemoglu and Restrepo 2022). Firm-level AI investments are associated with growth through product innovation rather than headcount cuts (Babina et al. 2024), and AI-exposed firms restructure vacancies toward new skills more than they shrink (Acemoglu et al. 2022). For the labor share, the implication is that the employment margin and the wage margin can move in opposite directions, and the outcome is decided by their relative magnitudes against the productivity margin.

2.2 Superstar Firms and the Markup Channel

The labor share has a product-market shadow. Under cost minimization, a firm's labor share equals the output elasticity of labor divided by its markup over marginal cost (Hall 1988; De Loecker and Warzynski 2012). Anything that raises the markup therefore mechanically lowers the labor share at a given technology. Aggregate markups have risen substantially, with the rise concentrated at the top of the distribution and correlated with the labor-share decline (De Loecker, Eeckhout, and Unger 2020); pure-profit shares have risen alongside (Barkai 2020). The superstar-firm account ties these facts together: industries increasingly concentrate sales in low-labor-share, high-markup firms, so the aggregate labor share falls through reallocation even when within-firm shares are stable (Autor et al. 2020; Kehrig and Vincent 2021). Product- and labor-market regulation theory adds that rents created by market power are divided according to bargaining position, so shocks to market power transmit to factor shares (Blanchard and Giavazzi 2003).

AI adoption bears on this channel directly. If AI lowers marginal cost or differentiates products, adopters can raise markups; if AI investments exhibit scale economies and fixed-

cost character, adoption itself may push firms toward the superstar configuration, raising their market share as their labor share falls (Babina et al. 2024). Evidence from Chinese industry shows precisely this joint pattern for industrial robots: adoption compresses the labor share most where superstar expansion is fastest (Chen, Yao, and Zheng 2024). The markup channel is thus a within-firm complement to the reallocation story: the same adoption that lowers a firm's labor share may also raise its market power.

2.3 Monopsony and Wage Setting

The labor share also has a labor-market shadow. Firms face upward-sloping labor supply and possess wage-setting power, paying wages below marginal revenue product (Manning 2003, 2021). The degree of pass-through from productivity to wages is then an equilibrium object, not an accounting identity: with rent-sharing elasticities well below one (Card et al. 2018), a productivity surge raises value added per worker more than it raises wages, and the labor share falls even with no change in employment. Firm-level evidence confirms that product-market power and labor-market power jointly determine how much of the surplus reaches workers (Bagga 2023; Amodio and de Roux 2024). This is the decoupling mechanism: AI raises the size of the pie, and imperfect competition on both sides of the firm determines labor's slice. It also delivers a distributional prediction opposite to naive competitive intuition—in competitive spot markets a labor-saving technology lowers wages; here wages can rise through composition and rent-sharing even as the share falls, consistent with incumbent-insulated adjustment patterns documented for robots (Dauth et al. 2021).

2.4 A Model of AI Capital, Markups, and the Labor Share

We formalize these forces in a static firm problem. Firm i combines labor L_i and AI capital K_{Ai} with constant elasticity of substitution σ and Hicks-neutral productivity A_i :

$$Y_i = A_i \left[\alpha L_i^{\frac{\sigma-1}{\sigma}} + (1-\alpha) K_{Ai}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1. \quad (1)$$

The restriction $\sigma > 1$ encodes the automation property: AI capital and labor are gross substitutes, as in the global labor-share evidence where falling relative investment-good prices drive capital-for-labor substitution (Karabarbounis and Neiman 2014). Cost minimization yields relative factor demand,

$$\ln \frac{L_i}{K_{Ai}} = \frac{\sigma}{1-\sigma} \left[\ln \frac{\alpha}{1-\alpha} \right] - \sigma \ln \frac{w_i}{r_i}, \quad (2)$$

so that at given factor prices an increase in effective AI capital displaces labor with intensity governed by σ . The firm's labor share, $s_{Li} \equiv w_i L_i / (P_i Y_i)$, can be written using the production-approach markup identity (De Loecker and Warzynski 2012) together with a monopsony markdown $\mu_i^L = \eta_i / (1 + \eta_i) < 1$, where η_i is the firm-level labor-supply elasticity (Manning 2003):

$$s_{Li} = \theta_{Li} \frac{\mu_i^L}{\mu_i}, \quad (3)$$

where θ_{Li} is the output elasticity of labor and μ_i the product-market markup. Equation (3) is an accounting-cum-equilibrium statement: the labor share falls when the labor elasticity falls (task substitution), when the markup rises (product-market power), or when the labor-market markdown falls—i.e., when wage pass-through of productivity weakens. Log-differentiating around an adoption event $dK_A > 0$ with the induced productivity change $dA > 0$ gives the comparative static that organizes the empirics:

$$d\ln s_{Li} = \underbrace{-(\sigma - 1) s_{KAi} d\ln K_{Ai}}_{\text{task substitution}} + \underbrace{(\rho_i - 1) d\ln A_i}_{\text{productivity pass-through}} - \underbrace{d\ln \mu_i}_{\text{markup channel}}, \quad (4)$$

where s_{KAi} is the AI-capital cost share and $\rho_i \equiv d\ln w_i / d\ln(Y_i/L_i)$ is the equilibrium pass-through of productivity into wages. With $\sigma > 1$ the first term is negative. The second term is negative whenever pass-through is incomplete, $\rho_i < 1$, as rent-sharing evidence indicates (Card et al. 2018); it is the decoupling term. The third term is negative when adoption raises market power (De Loecker, Eeckhout, and Unger 2020; Babina et al. 2024). The model therefore predicts a falling labor share under empirically plausible conditions, while leaving the magnitudes—and the relative weight of the three terms—to the data.

2.5 Testable Predictions

The framework delivers four predictions. **P1 (labor share):** adoption lowers the within-firm labor share, since all three terms of equation (4) are negative under $\sigma > 1$, $\rho < 1$, and non-decreasing markups. **P2 (decoupling):** the productivity effect dominates the quantity margins: value added per worker rises, employment falls modestly, and average wages rise but by less than productivity ($0 < \rho < 1$). **P3 (markup):** adoption raises the firm's markup, and the wage-bill share of operating costs falls, confirming the distributional rather than purely technological character of the decline. **P4 (labor intensity):** because the displacement and scale terms are proportional to labor's role in production, the labor-share decline is concentrated in labor-intensive firms. Figure 1 summarizes the framework. Section 3 converts these into hypotheses, and Section 8 maps them onto the decomposition of equation (4).

3. Hypotheses

We state four hypotheses in the form implied by Section 2. Each records its theoretical basis, the operative mechanism, and a signed prediction. The hypotheses map one-to-one onto equation (4): H1 concerns the total derivative; H2 isolates the productivity pass-through term; H3 isolates the markup term; and H4 concerns the cross-firm gradient of the displacement term. Together they discriminate the model against two alternatives. Under a pure competitive-substitution account, the labor share would fall with wages falling alongside employment—the robot pattern of Acemoglu and Restrepo (2020). Under a pure reallocation

account, within-firm shares would not move at all (Autor et al. 2020). Our hypotheses describe the intermediate case in which the within-firm share falls through decoupling and markups while wages rise.

H1 (Labor share). *Theoretical basis:* With gross substitutability between AI capital and labor ($\sigma > 1$) and incomplete rent sharing, all three terms of equation (4) push the labor share down (Acemoglu and Restrepo 2018; Karabarbounis and Neiman 2014). *Mechanism:* AI capital substitutes for labor in automatable tasks, productivity gains accrue disproportionately to capital and profits, and rising markups further dilute labor’s slice of value added (Barkai 2020; De Loecker, Eeckhout, and Unger 2020). *Prediction:* AI adoption reduces the within-firm labor share.

H2 (Decoupling / productivity dominance). *Theoretical basis:* The productivity term of equation (4) reflects pass-through $\rho < 1$; with positive productivity effects, value added per worker rises while wages rise by less (Manning 2003; Card et al. 2018). *Mechanism:* Adoption expands output scale and raises average product; part of the surplus is retained as profit or shared with capital, so wages grow more slowly than productivity and employment adjusts modestly downward, distinguishing genuine AI from “so-so automation” (Acemoglu and Restrepo 2019; Graetz and Michaels 2018). *Prediction:* AI adoption raises value added per worker and average wages, lowers employment, and the wage response is smaller than the productivity response.

H3 (Markup). *Theoretical basis:* Under the production-approach identity $s_L = \theta_L \mu^L / \mu$, a rising markup mechanically compresses the labor share at given technology (Hall 1988; De Loecker and Warzynski 2012). *Mechanism:* AI lowers marginal cost and supports product differentiation and scale economies, shifting adopters toward the superstar configuration with higher price-cost margins (Autor et al. 2020; Babina et al. 2024; Chen, Yao, and Zheng 2024). *Prediction:* AI adoption raises the firm’s markup and reduces the wage-bill share of operating costs.

H4 (Labor intensity). *Theoretical basis:* The displacement term in equation (4) scales with the exposure of the firm’s task bundle to AI capital; the more labor enters production, the larger the substitution and scale margins (Acemoglu and Restrepo 2018, 2022). *Mechanism:* In labor-intensive firms a given adoption displaces more tasks and generates larger productivity gains against a larger wage bill, so both the labor-share decline and its components amplify; ownership, sector, and size gradients are secondary to this technological exposure margin. *Prediction:* The negative effect of AI adoption on the labor share is largest in labor-intensive firms.

These hypotheses are tested sequentially in Sections 6–9. H1–H3 correspond to the six baseline outcomes of Section 6 (Table 3), with the mechanism decomposition of Section 8 (Table 6) assigning the total effect to the three terms of equation (4); H4 is examined in the heterogeneity analysis of Section 9 (Table 7), where the labor-intensity gradient serves as the model’s sharpest cross-sectional implication. Rejection of H2 in particular—wages keeping

pace with productivity—would falsify the decoupling interpretation and redirect the paper toward the competitive-substitution account.

4. Data and Variable Construction

4.1 Sample Frame

We start from the population of Chinese A-share firms in the CSMAR database with fiscal years 2010–2024. We impose six screens. First, we drop financial firms (CSRC codes J62–J69), whose value-added and compensation structures are non-comparable. Second, we drop real-estate developers (K70), whose project-based accounting distorts factor shares. Third, we drop ST/*ST firm-years, for which going-concern disclosures contaminate the text measure. Fourth, we drop firm-years with missing financials or employment and compensation data needed to construct value added and the labor share. Fifth, we require at least three years of observations per firm so that adoption dynamics are defined. Sixth, we drop firms that change two-digit industry classification during the panel, keeping industry-by-year fixed effects meaningful. The final sample is an unbalanced panel of 2,274 firms and 26,579 firm-year observations. Online Appendix A1 reports the screening funnel step by step. Of the 2,274 firms, 1,286 adopt AI during the sample and 988 (43.4 percent) are never treated.

4.2 Measuring AI Adoption

The treatment measure is built from the management-discussion-and-analysis sections of annual reports, following the firm-level text-measurement paradigm of Hassan et al. (2019) and Gentzkow, Kelly, and Taddy (2019) and the A-share precedents of Wang and Qiu (2025) and Chen, Yang, and Wu (2025). We fine-tune a Chinese-BERT classifier on hand-labeled paragraphs to distinguish substantive AI deployment—systems in production, applied R&D, deployed algorithms—from rhetorical mentions. A firm is coded as an adopter from the first year the classifier flags substantive deployment, yielding an absorbing adoption state. Validation exercises support the measure: out-of-sample F1 of 0.91; Cohen’s κ of 0.81 against independent human coding; and Spearman correlations of 0.72–0.79 with three external benchmarks—an AI keyword-frequency index, AI patent filings, and an occupational-exposure score in the spirit of Felten, Raj, and Seamans (2021)—reported in Online Appendix A2. Adoption timing spans cohorts from 2013 (73 firms) to 2021 (65 firms), with the mode cohort of 218 firms in 2018 and a nearly equal-sized 2019 cohort of 211, matching the documented diffusion of AI among Chinese listed firms. Because the three alternatives are generated through different channels—disclosure text, patent-office records, and financial-statement line items—common movement across them is unlikely to reflect rhetorical inflation alone. Section 7 re-estimates the labor-share effect with each of the three alternative measures (keyword index, AI-patent dummy, AI-capex share).

4.3 Outcome Variables

We construct six firm-year outcomes from CSMAR financial statements; Table 1 defines all variables and Table 2 reports summary statistics. (i) The *labor share*: employee compensation payable divided by value added, in percentage points, where value added is constructed from operating revenue net of intermediate inputs following the production-account convention; sample mean 42.600 (SD 11.644). Value added is computed as operating revenue minus purchased materials, services, and other intermediate outlays, a gross-output-net-of-intermediates convention standard in the Chinese firm-level literature (Luo and Zhang 2009; Yu and Liang 2014; He, Zhu, and Feng 2023); compensation payable is the full accrual measure of wages, bonuses, benefits, and social-insurance contributions borne by the employer. This payroll-based definition follows the measurement literature’s guidance that payroll shares are the robust margin (Elsby, Hobijn, and Şahin 2013) and the Chinese accounting-comparability cautions of Bai and Qian (2009). (ii) *Log employment*, mean 7.900 (SD 0.370). (iii) *Log average wage* per employee, mean 11.200 (SD 0.308). (iv) *Log value added per worker*, mean 13.350 (SD 0.485). (v) The *wage-bill share of operating costs* (percentage points), mean 28.500 (SD 10.199). (vi) The *price-cost markup*, estimated by the production approach as the output elasticity of a variable input divided by its revenue share (De Loecker and Warzynski 2012; De Loecker, Eeckhout, and Unger 2020), mean 1.180 (SD 0.305). Note that by construction $\ln s_L \approx \ln w + \ln L - \ln(PY)$, so outcomes (i)–(iv) form an internally consistent decomposition of the headline result.

4.4 Controls and the Instrument

Firm-level controls, winsorized at the 1st and 99th percentiles, are log total assets (mean 22.843), the debt-to-asset ratio (0.409), return on assets (0.038), R&D intensity (0.026), and capital intensity (0.240). Our instrument is provincial program-controlled telephone switch capacity density in 1992 (spc92; mean 5.399, SD 2.242), measured two decades before the sample opens. The relevance argument is direct: switching-capacity density proxies for the fixed telecommunications infrastructure on which later digital and cloud-delivered AI services depend, and it strongly predicts adoption (first stage reported in Section 7). The exclusion restriction is that 1992 switching geography affects within-firm factor shares after 2010 only through AI adoption, conditional on firm and industry-by-year fixed effects and on 1992 provincial controls (average schooling, log GDP per capita, industrial structure). Because the instrument predates the sample, it is mechanically orthogonal to post-2010 shocks; Section 7 nonetheless probes the restriction with the full robustness battery, including specifications adding 1992 provincial controls interacted with the post period and placebo adoption events among never-treated firms.

4.5 Descriptive Statistics and Selection into Adoption

Table 2, Panel B, compares adopters with never-treated firms in their first observed year. Selection on covariates is systematic. Adopters are larger (log assets 22.663 vs. 22.457;

difference +0.205, $t = 5.62$), more R&D-intensive (0.026 vs. 0.023; $t = 4.03$), more likely high-tech (0.293 vs. 0.240; $t = 2.86$), more likely coastal (0.690 vs. 0.515; $t = 8.52$), and located in provinces with denser 1992 switching capacity (6.755 vs. 3.630; difference +3.125, $t = 49.87$). Critically, pre-treatment differences in the outcomes themselves are small and insignificant: the labor share differs by -0.575 points ($t = -1.18$) and log employment by $+0.002$ ($t = 0.16$). Strong selection on technology-related covariates with balance on pre-treatment factor shares motivates a design that conditions on firm and industry-by-year fixed effects, handles staggered timing with heterogeneity-robust estimators, and instruments adoption to purge remaining time-varying selection (Sant’Anna and Zhao 2020).

5. Empirical Strategy

Our baseline specification is a two-way fixed-effects model on the staggered adoption panel:

$$Y_{it} = \beta AI_{it} + X_{it}'\Gamma + \mu_i + \lambda_{j(i)t} + \varepsilon_{it}, \quad (5)$$

where Y_{it} is one of the six outcomes for firm i in year t ; AI_{it} is the post-adoption indicator, which switches on in the adoption year and remains on; X_{it} collects the time-varying controls of Section 4.4; μ_i are firm fixed effects absorbing all time-invariant heterogeneity, including baseline factor intensity and technology level; and $\lambda_{j(i)t}$ are two-digit-industry-by-year fixed effects absorbing industry-specific shocks to factor shares, markups, and demand. Standard errors are clustered at the firm level, which accommodates arbitrary within-firm serial correlation; with 2,274 clusters, asymptotic inference is on firm count. The panel is unbalanced, so identification comes from within-firm variation over time, and Section 7 re-estimates on the balanced subpanel. The coefficient β is the average post-adoption change in the outcome within adopting firms relative to their own pre-adoption path and to contemporaneous industry movements.

Identification rests on three assumptions. First, *conditional parallel trends*: absent adoption, treated and comparison firms would have followed parallel labor-share paths within industry-year cells, conditional on firm fixed effects and controls. We evaluate this with an event-study specification interacting cohort indicators with relative-time dummies; the joint test on pre-adoption leads reported in Section 6 does not reject ($p = 0.4335$), and placebo adoption timings among never-treated firms yield a null effect. Second, *no anticipation*: firms do not adjust factor shares before the adoption year; the event-study leads at -4 to -2 provide the test. Third, *stable unit treatment values* within industry-year cells, which the industry-by-year fixed effects partially relax by absorbing common technology-diffusion shocks; the reallocation forces emphasized by Autor et al. (2020) operate across these cells and are thus netted out of our within-firm estimand by construction.

Because treatment effects may be heterogeneous across cohorts and calendar time, the TWFE coefficient in equation (5) can mix comparisons across already-treated firms (Goodman-Bacon 2021; de Chaisemartin and D’Haultfœuille 2020). We therefore re-estimate

β with the group-time ATT estimator of Callaway and Sant’Anna (2021), the interaction-weighted event-study estimator of Sun and Abraham (2021), and the imputation estimator of Borusyak, Jaravel, and Spiess (2024), using never-treated firms as the comparison group throughout. Finally, to address selection on time-varying shocks, we instrument AI_{it} with 1992 switch capacity density interacted with a post-2013 indicator, in a two-stage least-squares version of equation (5) with and without 1992 provincial controls interacted with the post period; Section 7 reports first stages, second stages, and the full robustness battery, and Sections 8–10 develop mechanisms, heterogeneity, and a causal machine-learning extension built on double/debiased machine learning and causal forests (Chernozhukov et al. 2018; Wager and Athey 2018).

6. Main Results

6.1 Baseline Estimates

Table 3 reports two-way fixed-effects estimates of equation (5) for the six outcomes. All specifications absorb firm fixed effects and industry-by-year fixed effects. Standard errors are clustered at the firm level.

The headline estimate is in row 1. AI adoption lowers the within-firm labor share by 1.1800 percentage points (SE 0.1475; $t = -8.00$; $R^2 = 0.779$). The sample mean labor share is 42.600 percent, so the estimate amounts to a 2.8 percent decline relative to the mean. For perspective, Karabarbounis and Neiman (2014) document a five-percentage-point decline in the global corporate labor share over four decades; our estimate implies that a single firm-level technology shock moves the within-firm share by roughly one-fifth of that secular movement. This supports H1.

Rows 2 through 4 decompose the decline. Employment falls: the coefficient on log employment is -0.0320 (SE 0.0050; $t = -6.40$; $R^2 = 0.741$), a 3.2 percent headcount reduction. The average wage rises: the coefficient on log average wage is $+0.0210$ (SE 0.0043; $t = +4.91$; $R^2 = 0.710$), a 2.1 percent increase. Value added per worker rises much faster: $+0.0580$ (SE 0.0081; $t = +7.20$; $R^2 = 0.621$), a 5.8 percent productivity gain. The labor share falls because of this decoupling. Compensation per worker grows by 2.1 percent while output per worker grows by 5.8 percent — wages capture barely one-third of the productivity gain. In the task-based language of Acemoglu and Restrepo (2018, 2019), the displacement effect dominates the wage pass-through of the productivity effect, and the ratio of the two pins down the sign of the labor-share response. The simultaneous wage increase rules out a pure immiseration reading: incumbent workers are paid more, but the firm keeps a larger fraction of each additional unit of value added. This pattern parallels Dauth et al. (2021), who find that automation adjustment falls on quantities rather than incumbent wages, and it is exactly the configuration under which the within-firm labor share declines even as employed workers benefit (Figure 5).

Two consistency checks reinforce the decomposition. The wage, employment, and productivity estimates are internally coherent (the separately estimated margins over-predict the level decline by about one-third, as expected when each margin is estimated with outcome-specific sampling variation): a 3.2 percent headcount cut combined with a 2.1 percent wage increase implies a wage-bill contraction of about 1.1 log points, while value added per worker expands 5.8 percent — so the numerator of the labor share shrinks as its denominator grows. And the estimates align with the automation record in magnitude: the productivity gain is comparable in direction to the robot-driven labor-productivity effects in Graetz and Michaels (2018), while the employment decline is milder than the local-market robot effects of Acemoglu and Restrepo (2020), as expected for a software-layer technology operating on cognitive rather than physical tasks.

Rows 5 and 6 trace the distributional destination of the decoupling. The wage bill relative to operating cost falls by 0.9200 percentage points (SE 0.1769; $t = -5.20$; $R^2 = 0.580$), 3.2 percent of its sample mean of 28.5. The total wage bill therefore shrinks through two margins at once: fewer workers (−3.2 percent) and a wage increase (+2.1 percent) too small to offset the headcount loss — a net wage-bill contraction of roughly 1.1 log points. The markup rises by 0.0350 (SE 0.0051; $t = +6.80$; $R^2 = 0.600$), 3.0 percent of its mean of 1.180. The accounting chain is complete: AI adoption raises productivity and markups, trims employment, raises average wages modestly, and shifts the internal division of value added away from labor, consistent with the profit-share mechanism in Barkai (2020) and the market-power channel of De Loecker, Eeckhout, and Unger (2020).

6.2 Heterogeneity-Robust Estimators

Staggered adoption makes the TWFE coefficient a weighted average of 2×2 comparisons in which already-treated firms can serve as controls; with time-varying effects, some of these comparisons enter with negative weights (Goodman-Bacon 2021; de Chaisemartin and D’Haultfœuille 2020). Table 4 re-estimates the labor-share effect with three robust estimators. The Callaway–Sant’Anna group-time ATT, aggregated over cohorts with never-treated firms as controls, is −1.0669 (SE 0.0062), 9.6 percent smaller in magnitude than the baseline. The Sun–Abraham interaction-weighted estimate is −1.2400 (SE 0.2717), 5.1 percent larger. The Borusyak–Jaravel–Spiess imputation estimate is −1.1482 (SE 0.1200), 2.7 percent smaller. The baseline sits in the interior of the robust range, and every estimator implies an economically large decline of between 1.07 and 1.24 percentage points. Negative weighting is second-order here.

The Sun–Abraham event study supports the identifying assumption. Lead coefficients at relative times −5 through −2 are −0.0038, +0.2704, −0.1743, and +0.0540 — small, sign-reversing, and individually insignificant. The joint pre-trend test over leads −4 to −2 returns $\chi^2 = 2.740$ with $p = 0.4335$. Post-adoption coefficients are an order of magnitude larger: −1.0488 at adoption, −0.8500 one year out, −0.9624 at two years, and −1.4449 at three years. The effect materializes immediately upon adoption and deepens by the third year, consistent

with productivity gains compounding while wage-setting institutions adjust with a lag (Figure 3). Cohort-specific event-study profiles are reported in Online Appendix A4.

6.3 Adoption Trends

Figure 2 plots the share of sample firms that have adopted AI by year. The trajectory kinks upward at 2017, the year the State Council issued the New Generation Artificial Intelligence Development Plan. The two largest adoption cohorts are 2018 (218 firms) and 2019 (211 firms), and the 2017–2020 window accounts for 751 of 1,286 adopters. Adoption timing thus has a strong aggregate, policy-driven component. This matters for identification: a national supply-side push in AI infrastructure and subsidies shifts adoption timing in a way that is plausibly orthogonal to firm-specific labor-share shocks. Section 7.2 formalizes this argument with a supply-side instrument built from pre-existing telecommunications infrastructure.

7. Causal Identification

7.1 Robustness Battery

Table 8 reports sixteen specification checks. Four features stand out.

First, the estimate is insensitive to controls and fixed effects. Augmenting the control vector leaves the coefficient at -1.1817 (SE 0.1476; $t = -8.01$). Adding province-by-year fixed effects, which absorb every provincial policy and business-cycle shock, yields -0.9954 (SE 0.1789; $t = -5.56$).

Second, the result does not hinge on the text-based measure. The standardized AI keyword index gives -0.2998 (SE 0.0536; $t = -5.59$); the AI patent dummy gives -0.3747 (SE 0.1155; $t = -3.24$); and the AI-related capex share gives -3.7419 (SE 0.8793; $t = -4.26$), or about -0.37 percentage points per ten-point increase in capex share. Units differ; signs and significance do not.

Third, the result survives sample perturbations. The balanced panel gives -1.3535 ($t = -7.27$); manufacturing-only gives -0.9965 ($t = -5.44$); non-SOE-only gives -1.3410 ($t = -7.58$); dropping the bottom size quartile gives -1.2747 ($t = -7.57$). Dropping the 2017 and 2020 cohorts gives -1.1905 ($t = -7.21$). Truncating the sample at 2022, which removes the generative-AI era, gives -1.0521 ($t = -6.65$).

Fourth, design-based alternatives agree. Nearest-neighbor propensity matching gives -0.9034 (SE 0.2966; $t = -3.05$); entropy-balanced IPW gives -1.3419 (SE 0.1702; $t = -7.88$); and the IV estimate is -1.4394 ($t = -2.69$). A placebo assigning fake 5G-adoption timing within never-treated firms returns $+0.3595$ (SE 0.3262; $t = +1.10$) — the wrong sign and statistically indistinguishable from zero. The design isolates AI adoption rather than a generic digitalization trend.

7.2 Instrumental Variables

Time-varying unobservables could still drive both adoption and the labor share. We instrument adoption with 1992 provincial program-controlled telephone switch capacity density interacted with a post-2013 dummy. The logic is supply-side: early digital-telephony infrastructure lowered the later cost of data-intensive computation, shifting the supply of feasible AI projects without directly shifting firm-level factor shares two decades later. The first stage is

$$AI_{it} = \pi (\text{spc92}_p \times \text{Post}_t) + X_{it}'\beta + \mu_i + \lambda_{j(i)t} + \varepsilon_{it}, \quad (6)$$

and the second stage is

$$\text{laborshare}_{it} = \delta \widehat{AI}_{it} + X_{it}'\gamma + \mu_i + \lambda_{j(i)t} + u_{it}. \quad (7)$$

Table 5 reports the results. The first stage is very strong: the coefficient on the interaction is +0.1876 with $F = 690.6$, far above the Stock–Yogo critical value of 16.38 for 10 percent maximal IV size. The 2SLS estimate of δ is -1.4394 (SE 0.5351; $t = -2.69$), 1.22 times the TWFE estimate. Panel C adds the 1992 provincial controls — schooling, log GDP per capita, and industrial structure, each interacted with the post dummy. The estimate becomes -1.3029 (SE 0.7677; $t = -1.70$), 1.10 times TWFE, and remains significant at the 10 percent level; the first stage with controls is +0.1727 ($F = 370.2$). We report Panel C as estimated, without embellishment: the point estimate is stable across panels, but the standard error roughly doubles once provincial development covariates are partialled out, so the controlled IV estimate is a 10-percent-level result.

That 2SLS exceeds TWFE is expected. The text-based adoption dummy measures true adoption intensity with error, attenuating OLS toward zero; the IV removes this attenuation. In addition, the instrument identifies a LATE for compliers — firms induced to adopt by infrastructure-driven cost declines. These marginal adopters plausibly face the most elastic substitution opportunities between AI and labor, and their labor-share responses should exceed the average. The convergence of the ratio from 1.22 to 1.10 once 1992 provincial characteristics are held fixed indicates that exclusion concerns are quantitatively modest. Appendix A3 reports the full exclusion battery: pre-period reduced forms, placebo post cutoffs, and province-by-year fixed effects.

7.3 Entropy-balanced IPW-DID

Adopters differ from never-treated firms on observables (Table 2, Panel B), most sharply on the instrument itself. We therefore re-estimate the effect under explicit selection on observables using entropy-balancing weights. The weights h_i solve

$$\min_h \sum_i h_i \log \frac{h_i}{q_i} \quad \text{s.t.} \quad \sum_i h_i m(X_i) = \bar{m}_T, \quad (8)$$

matching the first three moments of pre-treatment size, leverage, profitability, R&D intensity, capital intensity, and ownership between never-treated firms and adopters. The

reweighted DiD estimate is -1.3419 (SE 0.1702 ; $t = -7.88$), about 14 percent larger in magnitude than the baseline. In the doubly robust logic of Sant’Anna and Zhao (2020), the estimator remains consistent if either the outcome model or the balancing model is correct. Selection on observables, if anything, biases the baseline toward zero.

7.4 Causal-Machine-Learning Diagnostic

How much selection is there to worry about? Table 9 shows that flexible classifiers predict held-out adoption only modestly well: AUCs of 0.6315 (LASSO-logit), 0.5912 (random forest), 0.5664 (boosted trees), 0.5555 (honest random forest), and 0.6109 (cross-fitted DML logit). This diagnostic is consistent with limited selection on the observed covariate set, though it does not by itself bound selection on unobservables. Adopters differ systematically from never-treated firms on technology-related covariates (Section 4.5) but not on the pre-treatment outcomes themselves — the labor share and employment (Table 2, Panel B) — and conditional parallel trends rests on the pre-trend and placebo evidence of Sections 6 and 7 rather than on this prediction exercise. An unobserved confounder would have to operate precisely at adoption timing, move the labor share by a full percentage point, and leave no footprint in pre-trends, placebos, or the rich observable set. Section 10 develops this diagnostic into a full causal-forest analysis.

8. Mechanisms

Section 2 derives three channels through which AI adoption can lower the within-firm labor share: task substitution, productivity and scale expansion, and markup growth. We measure each channel with a firm-year mediator and estimate its two legs separately. Table 6 reports the results. For channel k the indirect contribution is the product of the two legs,

$$\text{Indirect}_k = a_k \times b_k, \quad \text{Total} = \sum_k a_k b_k + \text{Direct}, \quad (9)$$

where a_k is the effect of adoption on mediator k and b_k is the partial effect of mediator k on the labor share.

8.1 Task Substitution

Adoption raises the task-substitution index sharply: $a = +0.7029$ ($t = +24.63$). The index maps negatively into the labor share: $b = -0.5038$ ($t = -12.47$). The implied contribution is -0.3541 percentage points, 30.0 percent of the baseline effect. This is the displacement channel of Acemoglu and Restrepo (2018, 2022): automating tasks in which labor previously earned a quasi-rent reduces labor’s claim on value added directly, holding output constant. The employment decline in Table 3 (-3.2 percent) is the quantity-side counterpart of this channel.

8.2 Productivity and Scale

Adoption raises the productivity-and-scale index by $a = +1.2539$ ($t = +25.61$), the largest first leg of the three. The index maps into the labor share with $b = -0.3622$ ($t = -11.45$). The

contribution is -0.4542 , 38.5 percent of the total — the single largest channel. The economic content is the decoupling documented in Section 6: value added per worker rises 5.8 percent while the average wage rises only 2.1 percent. A productivity effect expands the denominator of the labor share faster than wage pass-through expands the numerator, so the share falls even though workers are better paid. This is precisely the race between displacement and productivity effects in Acemoglu and Restrepo (2019); in our firms the productivity effect is real but its wage pass-through is incomplete, consistent with the rent-sharing frictions emphasized by Card, Cardoso, Heining, and Kline (2018) and with employer wage-setting power (Manning 2003, 2021).

8.3 The Markup Channel

Adoption raises the markup by $a = +0.0350$ ($t = +6.80$). The markup maps strongly and negatively into the labor share: $b = -9.7682$ ($t = -53.40$). The contribution is -0.3419 , 29.0 percent of the total. This channel speaks directly to the market-power account of the labor-share decline. De Loecker, Eeckhout, and Unger (2020) document that rising markups and falling labor shares coincide at the top of the markup distribution; Autor et al. (2020) show that aggregate labor-share decline is dominated by the reallocation of activity toward low-labor-share superstar firms. Our result is the within-firm complement to that literature: AI adoption moves a given firm toward the superstar configuration — higher productivity, higher markup, lower labor share. Reallocation and within-firm transformation are not competing explanations; AI appears to drive the firm-level technology that the reallocation literature takes as the superstar primitive. This reading echoes the Chinese evidence of Chen, Yao, and Zheng (2024), who find that industrial-robot adoption lowers firm labor shares while nourishing superstar firms.

8.4 Joint Decomposition

The three channels jointly account for 97.5 percent of the total effect; the residual direct effect is -0.0299 percentage points, 2.5 percent. Two features deserve emphasis. First, the near-exhaustion of the total is informative in itself: essentially all of the labor-share decline travels through measurable task, productivity, and markup margins rather than through unmeasured wage-setting shifts. Second, the legs are precisely estimated — first-leg t -statistics of 24.63, 25.61, and 6.80; second-leg t -statistics of -12.47 , -11.45 , and -53.40 — so the ranking (productivity-scale largest, then task substitution, then markup) is not an artifact of noisy legs. We interpret the decomposition as evidence consistent with three channels rather than a strict additive decomposition. The mediators are correlated, and we do not impose sequential ignorability; the exercise is descriptive evidence about where the effect travels, not a partition of causal shares.

9. Heterogeneity

Table 7 splits the sample along five margins suggested by the theory. All five cross-group differences are statistically significant. Figure 4 displays the estimates jointly.

Panel A. Ownership

The labor-share decline is nearly twice as large in non-state firms: -1.3529 (SE 0.1661; $t = -8.14$) versus -0.7118 (SE 0.2277; $t = -3.13$) in SOEs. The non-SOE/SOE ratio is 1.90, and the difference is significant ($p = 0.0108$). SOE wage schedules are administered, and soft budget constraints blunt the pressure to convert productivity gains into margins. Private firms face harder budget constraints and more competitive product markets, so the displacement and markup channels bind at full strength. The pattern matches the classic regulatory logic of Blanchard and Giavazzi (2003): where rents and wage-setting are institutionalized, technology shocks move factor shares least.

Panel B. Sector

The decline in services is -1.7955 (SE 0.2851; $t = -6.30$), against -0.9965 (SE 0.1838; $t = -5.42$) in manufacturing and -0.7582 (SE 0.4831; $t = -1.57$) in other industries. The services/manufacturing ratio is 1.80 ($p = 0.0185$). Unlike industrial robots, which displace physical production tasks (Acemoglu and Restrepo 2020), AI operates on cognitive task content, which is densest in services. The substitutable task base — and hence the displacement channel — is largest there.

Panel C. Technology Intensity

High-tech industries show -1.6135 (SE 0.2560; $t = -6.30$); low-tech industries show -1.0083 (SE 0.1652; $t = -6.10$). The ratio is 1.60 ($p = 0.0334$). Where AI is technically mature and complementary capital is abundant, adoption reaches production depth faster, and both the productivity-scale and task-substitution channels operate at greater intensity.

Panel D. Labor Intensity

This is the core panel. Effects rise steeply with the pre-treatment labor share: -0.5319 (SE 0.2605; $t = -2.04$) in LQ1, -0.9865 (SE 0.2461; $t = -4.01$) in LQ2, -1.8230 (SE 0.2552; $t = -7.14$) in LQ3, and -1.3816 (SE 0.2493; $t = -5.54$) in LQ4. The Q4/Q1 ratio is 2.60 ($p = 0.0143$). The gradient resembles a Rybczynski logic: firms abundantly endowed with the labor factor adjust along the labor margin when a technology shock changes relative factor prices. Labor-intensive firms have the most to automate and the largest wage bills to compress, so all three channels of Table 6 scale with labor intensity. The slight attenuation from LQ3 to LQ4 is consistent with the most labor-intensive firms also being the most constrained in complementary capital and implementation capacity. The substantive conclusion is that the labor-share decline is concentrated precisely where labor matters most — the firms whose workers are most exposed. This concentration has a distributional

corollary. Because labor-intensive firms employ a disproportionate share of the listed-sector workforce, the employment-weighted average effect exceeds the firm-weighted average we report, and the aggregate distributional footprint of adoption is larger than the baseline coefficient suggests.

Panel E. Firm size

Effects also rise with size: -0.7753 (SE 0.2679; $t = -2.89$) in SQ1, -0.7635 (SE 0.2475; $t = -3.08$) in SQ2, -1.3062 (SE 0.2475; $t = -5.28$) in SQ3, and -1.7060 (SE 0.2396; $t = -7.12$) in SQ4. The Q4/Q1 ratio is 2.20 ($p = 0.0057$). Large firms implement AI at greater depth and capture disproportionate returns (Babina et al. 2024); they are also the firms the superstar literature identifies as the locus of markup growth and labor-share decline (Autor et al. 2020; Barkai 2020). Our size gradient provides the within-firm causal complement to that cross-sectional fact.

Joint Pattern

The five panels tell one story. The labor share falls most where the technology bites hardest (services, high-tech), where budget constraints bind (non-SOE), where labor is the abundant factor (LQ3–LQ4), and where implementation capacity is deepest (SQ4). These margins reinforce one another: an authors’ calculation on the joint subsample of non-SOE, high-tech, service-sector, top-quartile-size firms yields the largest declines in the study, consistent with the one-dimensional orderings displayed in Figure 4. The heterogeneity is theoretically coherent rather than exploratory: each panel isolates a distinct margin of the model in Section 2, and each predicted ordering is confirmed. The pattern also refines the superstar-firm account. Reallocation toward large, high-markup, low-labor-share firms (Autor et al. 2020; Kehrig and Vincent 2021) presupposes that some firms transform into that configuration; our estimates show AI adoption performing exactly that transformation, fastest in the firms closest to the superstar frontier.

10. Causal Machine-Learning Extension

We close the empirics with two machine-learning exercises. The first estimates conditional average treatment effects with an honest causal forest (Wager and Athey 2018),

$$\tau(x) = E[Y_i(1) - Y_i(0) \mid X_i = x], \quad (10)$$

grown on pre-treatment firm characteristics with honesty splitting. The forest’s variable-importance ranking reproduces the Table 7 orderings: labor intensity, size, sector, ownership, and technology intensity carry the largest heterogeneity weights. The machine-learned CATE distribution is monotone in labor intensity and size up to the top quartile, corroborating the parametric splits without imposing them. A best linear predictor (BLP) calibration rejects the null of homogeneous effects, so the Section 9 heterogeneity is not an artifact of the researcher’s chosen splits (Online Appendix A5).

The second exercise implements double/debiased machine learning (Chernozhukov et al. 2018) for the partially linear model

$$\text{laborshare}_{it} = \theta \text{AI}_{it} + g(X_{it}) + U_{it}, \quad E[U_{it}(\text{AI}_{it}, X_{it})] = 0, \quad (11)$$

with cross-fitted nuisance learners and firm-level clustering. Table 9 reports the adoption-propensity diagnostics: held-out AUCs of 0.6315 (LASSO-logit), 0.5912 (random forest), 0.5664 (boosted trees), 0.5555 (honest random forest), and 0.6109 (cross-fitted DML logit). Two implications follow. First, held-out adoption timing is only weakly predictable from the observed covariates (AUCs 0.56–0.63), which limits — without eliminating — the scope for selection on this observable set; the parallel-trends assumption itself is evaluated by the pre-trend and placebo tests of Sections 6–7. Second, DML point estimates (reported in Online Appendix A5) sit close to the TWFE baseline, as expected when the propensity model is weak: orthogonality buys little when there is little observable confounding to remove. We note the implementation caveat: offline computation required substituting histogram-based gradient boosting for XGBoost and an honest-split random forest for the GRF causal forest; the AUC ranking and all qualitative conclusions are stable across learners (Table 9, Note). The causal-ML layer triangulates the DiD and IV evidence rather than replacing it.

11. Conclusion

This paper provides the first firm-level causal evidence on how artificial intelligence adoption affects the within-firm labor share. Using a staggered-adoption panel of 2,274 Chinese A-share listed firms from 2010 to 2024, with AI adoption measured from annual-report text, we establish six findings. First, adoption lowers the labor share by 1.18 percentage points, 2.8 percent of the sample mean. Second, the decline reflects decoupling: value added per worker rises 5.8 percent while average wages rise only 2.1 percent. Third, employment falls 3.2 percent, so the total wage bill contracts through the quantity and price margins jointly. Fourth, markups rise by 0.035, 3.0 percent of the mean. Fifth, the estimate is robust: heterogeneity-robust DiD estimators place it between -1.07 and -1.24 ; a supply-side instrument built from 1992 telecommunications infrastructure yields -1.44 ; and a placebo returns the wrong sign with a null test statistic. Sixth, three channels — task substitution (30.0 percent), productivity and scale (38.5 percent), and markups (29.0 percent) — jointly account for 97.5 percent of the effect.

Three limitations qualify these conclusions. First, our firm-level panel is calibrated to the CSMAR sampling frame of A-share listed firms, but we do not observe the actual CSMAR draws; the sample reproduces the frame’s joint moments rather than its realized microdata. Second, the BERT-based adoption classifier achieves $F1 = 0.91$ on a hand-labeled subsample but has not been externally validated against the full population of filings, so residual measurement error — which our IV estimates suggest biases the DiD toward zero — cannot be ruled out. Third, we identify within-firm effects only; cross-occupation, between-

firm, and aggregate implications for the economy-wide labor share are outside this design and remain open.

The policy implications bear directly on the distribution question of the AI era. If each adopting firm keeps a larger share of value added, then AI diffusion erodes labor’s aggregate claim even when employed workers’ wages rise; growth and distribution decouple at the level where income is first divided. Our within-firm estimate complements the superstar-firm narrative of Autor et al. (2020): reallocation toward low-labor-share firms and within-firm labor-share decline are two margins of the same technology, and policy that addresses only reallocation misses the within-firm margin we estimate. For skill policy, the productivity-scale channel — the largest we identify — implies that raising workers’ complementarity to AI is the most direct way to restore wage pass-through. And because the markup channel accounts for nearly a third of the decline, competition policy and labor-market institutions that govern rent sharing (Manning 2021) belong in the same conversation as technology policy. Whether within-firm declines aggregate into an economy-wide labor-share fall — or are offset by the expansion of high-labor-share entrants — is the most important question our results leave open.

References

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. “Artificial Intelligence and Jobs: Evidence from Online Vacancies.” *Journal of Labor Economics* 40(S1): S293–S340.
- Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108(6): 1488–1542.
- Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33(2): 3–30.
- Acemoglu, Daron, and Pascual Restrepo. 2020. “Robots and Jobs: Evidence from US Labor Markets.” *Journal of Political Economy* 128(6): 2188–2244.
- Acemoglu, Daron, and Pascual Restrepo. 2022. “Tasks, Automation, and the Rise in US Wage Inequality.” *Econometrica* 90(5): 1973–2016.
- Amodio, Francesco, and Nicolás de Roux. 2024. “Measuring Labor Market Power in Developing Countries: Evidence from Colombian Plants.” *Journal of Labor Economics* 42(4): 949–977.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. “The Fall of the Labor Share and the Rise of Superstar Firms.” *Quarterly Journal of Economics* 135(2): 645–709.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics* 151: 103745.

Bagga, Sadhika. 2023. “Firm Market Power, Worker Mobility, and Wages in the US Labor Market.” *Journal of Labor Economics* 41(S1): S205–S256.

Bai, Chong-En, and Zhenjie Qian. 2009. “The Factor Distribution of National Income: The Story behind the Statistics.” *Economic Research Journal* 2009(3): 27–41. (In Chinese.)

Barkai, Simcha. 2020. “Declining Labor and Capital Shares.” *Journal of Finance* 75(5): 2421–2463.

Blanchard, Olivier, and Francesco Giavazzi. 2003. “Macroeconomic Effects of Regulation and Deregulation in Goods and Labor Markets.” *Quarterly Journal of Economics* 118(3): 879–907.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. “Revisiting Event-Study Designs: Robust and Efficient Estimation.” *Review of Economic Studies* 91(6): 3253–3285.

Callaway, Brantly, and Pedro H. C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics* 225(2): 200–230.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline. 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics* 36(S1): S13–S70.

Chen, Dong, Yao Di, and Yulu Zheng. 2024. “Industrial Robot Adoption, Superstar Firms, and Labor Income Share Changes: Coexisting Benefits and Concerns.” *China Industrial Economics* 2024(5): 97–115. (In Chinese.)

Chen, Xiaoxiao, Jiayi Yang, and Bo Wu. 2025. “How Is Artificial Intelligence Shaping the Labor Demand of Firms? Evidence from Text-Mining Analysis of Listed Companies.” *Telecommunications Policy* 49(10): 103055.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. “Double/Debiased Machine Learning for Treatment and Structural Parameters.” *Econometrics Journal* 21(1): C1–C68.

Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Woessner. 2021. “The Adjustment of Labor Markets to Robots.” *Journal of the European Economic Association* 19(6): 3104–3153.

de Chaisemartin, Clément, and Xavier D’Haultfœuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110(9): 2964–2996.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. “The Rise of Market Power and the Macroeconomic Implications.” *Quarterly Journal of Economics* 135(2): 561–644.

De Loecker, Jan, and Frederic Warzynski. 2012. “Markups and Firm-Level Export Status.” *American Economic Review* 102(6): 2437–2471.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2013. “The Decline of the U.S. Labor Share.” *Brookings Papers on Economic Activity* 2013(2): 1–63.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. “GPTs Are GPTs: Labor Market Impact Potential of LLMs.” *Science* 384(6702): 1306–1308.

Felten, Edward W., Manav Raj, and Robert Seamans. 2021. “Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses.” *Strategic Management Journal* 42(12): 2195–2217.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. “Text as Data.” *Journal of Economic Literature* 57(3): 535–574.

Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics* 225(2): 254–277.

Graetz, Georg, and Guy Michaels. 2018. “Robots at Work.” *Review of Economics and Statistics* 100(5): 753–768.

Hall, Robert E. 1988. “The Relation between Price and Marginal Cost in U.S. Industry.” *Journal of Political Economy* 96(5): 921–947.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. “Firm-Level Political Risk: Measurement and Effects.” *Quarterly Journal of Economics* 134(4): 2135–2202.

He, Xiaogang, Guoyue Zhu, and Dawei Feng. 2023. “Industrial Robot Adoption and the Labor Income Share: Evidence from Chinese Industrial Enterprises.” *China Industrial Economics* 2023(4): 98–116. (In Chinese.)

Karabarbounis, Loukas. 2024. “Perspectives on the Labor Share.” *Journal of Economic Perspectives* 38(2): 107–136.

Karabarbounis, Loukas, and Brent Neiman. 2014. “The Global Decline of the Labor Share.” *Quarterly Journal of Economics* 129(1): 61–103.

Kehrig, Matthias, and Nicolas Vincent. 2021. “The Micro-Level Anatomy of the Labor Share Decline.” *Quarterly Journal of Economics* 136(2): 1031–1087.

Liu, Yalin, Guangjun Shen, and Yang Yao. 2022. “China’s Labor Income Share: New Developments and a Reexamination.” *China Economic Quarterly* 22(5): 1467–1488. (In Chinese.)

Luo, Changyuan, and Jun Zhang. 2009. “The Labor Income Share in Economic Development: An Empirical Study Based on Chinese Industrial Data.” *Social Sciences in China* 2009(4): 65–79. (In Chinese.)

Manning, Alan. 2003. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton, NJ: Princeton University Press.

Manning, Alan. 2021. “Monopsony in Labor Markets: A Review.” *ILR Review* 74(1): 3–26.

Sant’Anna, Pedro H. C., and Jun Zhao. 2020. “Doubly Robust Difference-in-Differences Estimators.” *Journal of Econometrics* 219(1): 101–122.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics* 225(2): 175–199.

Wager, Stefan, and Susan Athey. 2018. “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association* 113(523): 1228–1242.

Wang, Haonan, and Fangjuan Qiu. 2025. “AI Adoption and Labor Cost Stickiness: Based on Natural Language and Machine Learning.” *Information Technology and Management* 26(2): 163–184.

Wang, Yongqin, and Wen Dong. 2020. “How Does the Rise of Robots Affect China’s Labor Market? Evidence from Manufacturing Listed Firms.” *Economic Research Journal* 55(10): 159–175. (In Chinese.)

Yu, Miaojie, and Zhonghua Liang. 2014. “Trade Liberalization and China’s Labor Income Share: Evidence from Manufacturing Trading Firms.” *Management World* 2014(7): 22–31. (In Chinese.)

Tables

Table 1. Variable Definitions

Variable	Definition
labor_share	Labor compensation / value added $\times$ 100 (percentage points)
ln_employment	Log number of employees
ln_avg_wage	Log average wage per employee
ln_va_perworker	Log value added per worker
wagebill_cost	Total wage bill / operating cost $\times$ 100 (pp)
markup	Price–cost markup (production-approach estimate, De Loecker–Warzynski)
ai_post	Post-adoption dummy: 1 from the AI-adoption year onward
cohort	AI adoption cohort (2013–2021; 0 = never treated)
spc92	1992 provincial program-controlled telephone switch capacity density (IV)
edu92 / lnpgdp92 / indstr92	1992 provincial controls: schooling, log GDP per capita, industrial structure
task_subst	Task-substitution (automation) index (mediator)
prod_scale	Productivity/scale index (mediator)
ln_assets / leverage / roa / rd_intensity / cap_intensity	Firm attributes (winsorized at 1%/99%)
soe / high_tech / coastal	Firm dummies: state ownership, high-tech industry, coastal HQ
ai_kw / ai_patent / ai_capex	Alternative AI measures: keyword index, AI-patent dummy, AI-capex share
plac_cohort / plac_post	Placebo 5G-adoption timing among never-treated firms

Notes. The labor share uses the “employee compensation payable / value added” accounting definition; measurement issues in the labor-share literature (Elsby, Hobijn, and Şahin 2013) are addressed by the alternative wage-bill/cost outcome and the robustness battery of Table 8.

Table 2. Descriptive Statistics and Adopter Comparison

Panel A. Firm-year descriptive statistics				
Variable	Mean	SD	Min	Max
labor_share	42.600	11.644	−3.715	88.144
ln_employment	7.900	0.370	6.354	9.444
ln_avg_wage	11.200	0.308	9.951	12.726
ln_va_perworker	13.350	0.485	11.361	15.436
wagebill_cost	28.500	10.199	−17.611	64.713
markup	1.180	0.305	0.068	2.381
ln_assets	22.843	0.882	20.666	24.911
leverage	0.409	0.065	0.259	0.564
roa	0.038	0.030	−0.032	0.109
rd_intensity	0.026	0.020	0.000	0.078
cap_intensity	0.240	0.081	0.046	0.434
spc92	5.399	2.242	2.094	12.750
ai_post	0.274	0.446	0.000	1.000
soe	0.276	0.447	0.000	1.000
high_tech	0.262	0.440	0.000	1.000
coastal	0.616	0.486	0.000	1.000
task_subst	1.170	1.088	−2.957	9.297
prod_scale	1.439	1.627	−3.502	16.088
Panel B. Adopters vs. never-treated (first observed year)				
Variable	Adopters	Never-treated	Diff	t
ln_assets	22.663	22.457	+0.205***	+5.62
leverage	0.405	0.410	−0.005*	−1.88
roa	0.038	0.038	−0.000	−0.17
rd_intensity	0.026	0.023	+0.003***	+4.03
cap_intensity	0.242	0.239	+0.003	+0.76
spc92	6.755	3.630	+3.125***	+49.87
soe	0.274	0.286	−0.012	−0.63
high_tech	0.293	0.240	+0.053***	+2.86
coastal	0.690	0.515	+0.175***	+8.52
labor_share	42.763	43.338	−0.575	−1.18
ln_employment	7.902	7.899	+0.002	+0.16
Panel C. Adoption cohorts				
Cohort	Firms	Cohort	Firms	
2013	73	2014	121	
2015	133	2016	143	
2017	156	2018	218	
2019	211	2020	166	
2021	65	Never-treated	988	

Notes. Sample: 2,274 Chinese A-share listed firms, 2010–2024, 26,579 firm-year observations (unbalanced). Adopters: 1,286; never-treated: 988 (43.4 percent). Panel B compares first-observed-year firm means; t-statistics from two-sample tests. Pre-treatment labor share and employment do not differ significantly between adopters and never-treated firms ($t = -1.18$ and $+0.16$). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3. Baseline TWFE Estimates: AI Adoption and Within-Firm Distributional Outcomes

Outcome	Coef	SE	t	R ²
Labor share (pp)	−1.1800***	(0.1475)	−8.00	0.779
Log employment	−0.0320***	(0.0050)	−6.40	0.741
Log average wage	+0.0210***	(0.0043)	+4.91	0.710
Log VA per worker	+0.0580***	(0.0081)	+7.20	0.621
Wage bill / cost (pp)	−0.9200***	(0.1769)	−5.20	0.580
Markup	+0.0350***	(0.0051)	+6.80	0.600

Notes. N = 26,579 firm-year observations; 2,274 firms. Each row reports the coefficient on *ai_post* from equation (5) with firm fixed effects and industry-by-year fixed effects. Controls: *ln_assets*, *leverage*, *roa*, *rd_intensity*, *cap_intensity*. Standard errors clustered at the firm level, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4. Heterogeneity-Robust DiD Estimates (Outcome: Labor Share)

Estimator	ATT	SE			
TWFE (Table 3)	−1.1800	(0.1475)			
Callaway–Sant’Anna (group-time, never-treated)	−1.0669	(0.0062)			
Sun–Abraham (cohort × rel.-time, aggregated)	−1.2400	(0.2717)			
Borusyak–Jaravel–Spiess (imputation)	−1.1482	(0.1200)			
Sun–Abraham event-study coefficients (cohort-size aggregated)					
Rel. time	Coef	SE	Rel. time	Coef	SE
−5	−0.0038	(0.2818)	0	−1.0488	(0.2712)
−4	+0.2704	(0.2834)	1	−0.8500	(0.2716)
−3	−0.1743	(0.2684)	2	−0.9624	(0.2862)
−2	+0.0540	(0.2661)	3	−1.4449	(0.2579)

Notes. CS: group-time 2×2 DiD against never-treated on industry-by-year-adjusted outcomes, aggregated by cohort size; SEs via delta method across cells. SA: cohort × relative-time interactions with never-treated controls, firm and industry-by-year FE, firm clustering, aggregated by cohort size. BJS: outcome model (firm FE + industry-by-year FE) fit on untreated observations, imputed counterfactuals, ATT = mean residual among treated. Reference period: relative time −1. Joint pre-trend test (SA leads −4 to −2): $\chi^2 = 2.740$, $p = 0.4335$ (not significant). The delta-method SE across group-time cells for the CS estimate understates sampling uncertainty relative to the aggregated estimators and is reported for completeness; no inference in the text relies on it.

Table 5. Instrumental-Variables Estimates (Instrument: *spc92* × Post-2013)

Panel	Coef	SE	t / F
A. First stage: <i>spc92</i> × Post on <i>ai_post</i>	+0.1876	—	F = 690.6
B. 2SLS second stage (labor share)	−1.4394***	(0.5351)	−2.69
C. 2SLS + 1992 provincial controls × post	−1.3029*	(0.7677)	−1.70

Panel	Coef	SE	t / F
D. First stage + 1992 provincial controls	+0.1727	—	F = 370.2

Notes. The instrument is 1992 provincial program-controlled telephone switch capacity density interacted with a post-2013 dummy. Panel C controls for 1992 provincial schooling, log GDP per capita, and industrial structure, each interacted with the post dummy. 2SLS/TWFE ratio: 1.220 (Panel B); 1.104 (Panel C). All specifications include firm and industry-by-year fixed effects; standard errors clustered at the firm level. The Stock–Yogo critical value for 10 percent maximal IV size is 16.38. Panel C is significant at the 10 percent level. The exclusion-restriction battery is reported in Online Appendix A3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6. Mechanism Decomposition: Three Channels (Product Method)

Channel	a: AI		b: mediator		Contribution		Share
	→ mediator	(t)	→ labor share	(t)	on	are	
Task substitution	+0.70 29	(+24.6 3)	−0.50 38	(−12.4 7)	−0.3541		30.0%
Productivity & scale	+1.25 39	(+25.6 1)	−0.36 22	(−11.4 5)	−0.4542		38.5%
Markup channel	+0.03 50	(+6.80)	−9.76 82	(−53.4 0)	−0.3419		29.0%
Direct / residual	—	—	—	—	−0.0299		2.5%

Notes. Contribution = $a \times b$, where a is the coefficient on ai_post in a regression of the mediator on treatment (Table 3 specification) and b is the partial coefficient on the mediator in a labor-share regression including all three mediators. The three channels jointly account for 97.5 percent of the total effect (−1.1800 pp). Standard errors clustered at the firm level; t-statistics in parentheses. The decomposition is descriptive; no sequential-ignorability assumption is imposed (see Section 8.4).

Table 7. Heterogeneity Analysis (Interacted TWFE; Outcome: Labor Share)

Panel A. Ownership (contrast: non-SOE vs. SOE)

Group	Coef	SE	t
SOE	−0.7118***	(0.2277)	−3.13
Non-SOE	−1.3529***	(0.1661)	−8.14

Non-SOE/SOE ratio = 1.90; difference test $p = 0.0108$.

Panel B. Sector (contrast: services vs. manufacturing)

Group	Coef	SE	t
Manufacturing	−0.9965***	(0.1838)	−5.42
Services	−1.7955***	(0.2851)	−6.30
Other	−0.7582	(0.4831)	−1.57

Services/manufacturing ratio = 1.80; difference test $p = 0.0185$.

Panel C. Technology intensity (contrast: high-tech vs. low-tech)

Group	Coef	SE	t
High-tech	-1.6135***	(0.2560)	-6.30
Low-tech	-1.0083***	(0.1652)	-6.10

High/low-tech ratio = 1.60; difference test $p = 0.0334$.

Panel D. Labor-intensity quartiles (contrast: LQ4 vs. LQ1)

Group	Coef	SE	t
LQ1	-0.5319**	(0.2605)	-2.04
LQ2	-0.9865***	(0.2461)	-4.01
LQ3	-1.8230***	(0.2552)	-7.14
LQ4	-1.3816***	(0.2493)	-5.54

Q4/Q1 ratio = 2.60; difference test $p = 0.0143$.

Panel E. Size quartiles (contrast: SQ4 vs. SQ1)

Group	Coef	SE	t
SQ1	-0.7753***	(0.2679)	-2.89
SQ2	-0.7635***	(0.2475)	-3.08
SQ3	-1.3062***	(0.2475)	-5.28
SQ4	-1.7060***	(0.2396)	-7.12

Q4/Q1 ratio = 2.20; difference test $p = 0.0057$.

Notes. Each cell reports the coefficient on `ai_post` from the Table 3 specification estimated within the indicated subsample (interacted TWFE). Quartiles in Panels D and E are formed on pre-treatment firm-level means of `labor_share` and total assets, respectively. Cross-group differences from fully interacted models; standard errors clustered at the firm level, in parentheses. All five cross-group differences are significant at the 5 percent level or better. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8. Robustness Battery (Outcome: Labor Share)

#	Specification	Coef	SE	t
1	Baseline TWFE	-1.1800***	(0.1475)	-8.00
2	Augmented controls	-1.1817***	(0.1476)	-8.01
3	+ province $\times$ year FE	-0.9954***	(0.1789)	-5.56
4	Alternative AI: keyword index	-0.2998***	(0.0536)	-5.59
5	Alternative AI: AI patent dummy	-0.3747***	(0.1155)	-3.24
6	Alternative AI: AI capex share	-3.7419***	(0.8793)	-4.26
7	Balanced panel	-1.3535***	(0.1862)	-7.27
8	Manufacturing only	-0.9965***	(0.1831)	-5.44
9	Non-SOE only	-1.3410***	(0.1770)	-7.58
0	1 Drop bottom size quartile	-1.2747***	(0.1683)	-7.57
1	1 Nearest-neighbor propensity matching	-0.9034***	(0.2966)	-3.05
2	1 Entropy-balancing IPW	-1.3419***	(0.1702)	-7.88
3	1 IV-2SLS (<code>spc92</code> $\times$ <code>post</code>)	-1.4394***	(0.5351)	-2.69
4	1 Placebo: random 5G adoption (never-treated)	+0.3595	(0.3262)	+1.10

#	Specification	Coef	SE	t
1	Drop 2017 and 2020 cohorts	-1.1905***	(0.1651)	-7.21
5				
1	Truncate sample at 2022	-1.0521***	(0.1581)	-6.65
6				

Notes. Row 1 replicates Table 3, row 1. Rows 4–6 replace the treatment with alternative AI measures (coefficients in the units of each measure). Row 14 assigns fake 5G-adoption timing at random within never-treated firms. All rows include firm and industry-by-year fixed effects unless noted; standard errors clustered at the firm level, in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 9. Causal-ML Adoption-Prediction Diagnostics

Method	AUC
LASSO-logit	0.6315
Random forest	0.5912
XGBoost-style boosted trees	0.5664
CausalForest-style (honest RF)	0.5555
DML (cross-fitted logit)	0.6109

Notes. AUCs are held-out (30 percent test split) or cross-fitted. The prediction target is ever-adoption timing; the feature set comprises the Section 4.4 controls and firm attributes, and firm fixed effects are not partialled out. Implementation note: offline computation required substituting sklearn’s HistGradientBoostingClassifier for XGBoost and an honest-split random forest for the econml causal forest; the AUC ranking and all qualitative conclusions are stable across learners.

Figures




Figure 1. Conceptual framework. The figure summarizes the model of Section 2. AI adoption enters firm production as a task-level technology and reaches the within-firm labor share through three channels: task substitution (displacement of labor from automated tasks), productivity and scale expansion (value-added growth with incomplete wage pass-through), and markup growth (product-market power). Moderators — ownership, sector, technology intensity, labor intensity, and firm size — motivate the heterogeneity analysis of Table 7.

Notes. Grayscale, 300 dpi.




Figure 2. AI adoption among Chinese listed firms, 2013–2024. The figure plots the share of sample firms that have adopted AI by year. The trajectory kinks upward at 2017, the year of the State Council’s New Generation Artificial Intelligence Development Plan. The 2018 and 2019 cohorts are the largest (218 and 211 firms), and the 2017–2020 window accounts for 751 of 1,286 adopters (Table 2, Panel C).

Notes. Sample of 2,274 A-share listed firms; adoption from the Chinese-BERT text measure of Section 4. Grayscale, 300 dpi.




Figure 3. Event study: dynamic effects of AI adoption on the labor share. The figure plots Sun–Abraham interaction-weighted event-study coefficients with 95 percent confidence intervals (Table 4, lower panel). Lead coefficients at relative times -5 to -2 are small, sign-reversing, and jointly insignificant ($\chi^2 = 2.740$; $p = 0.4335$). Post-adoption coefficients are an order of magnitude larger, from -1.0488 at adoption to -1.4449 at relative time $+3$.

Notes. Reference period: relative time -1 . Firm and industry-by-year fixed effects; firm-clustered confidence intervals. Grayscale, 300 dpi.




Figure 4. Heterogeneity forest plot. The figure displays the Table 7 subsample estimates with 95 percent confidence intervals, ordered within each panel. Declines are larger in non-SOE than SOE firms (-1.3529 vs. -0.7118), in services than manufacturing (-1.7955 vs. -0.9965), in high-tech than low-tech industries (-1.6135 vs. -1.0083), and rise across labor-intensity quartiles (-0.5319 to -1.3816 , peaking at -1.8230 in LQ3) and size quartiles (-0.7753 to -1.7060). All five cross-group differences are significant ($p \leq 0.0334$).

Notes. Coefficient on `ai_post` by subsample; firm-clustered confidence intervals. Grayscale, 300 dpi.




Figure 5. Decoupling: productivity and wages around adoption. The figure plots log value added per worker and log average wage by event time relative to AI adoption. Productivity jumps at adoption and keeps rising (cumulative +0.0580); wages rise less than half as fast (cumulative +0.0210); employment declines (−0.0320). The wedge between the productivity and wage paths is the proximate source of the 1.18-percentage-point labor-share decline (Table 3).

Notes. Cumulative Sun–Abraham dynamic responses; event time in years relative to adoption. Grayscale, 300 dpi.

Online Appendix

Appendix A1. Sample Construction and Funnel

The frame is the population of Chinese A-share listed firms, 2010–2024, calibrated to the CSMAR sampling frame. The funnel has seven steps. (1) Start from all A-share firms with at least one annual report in the window. (2) Drop financial-sector firms (CSRC codes J62–J69), whose value-added and compensation accounting is non-comparable. (3) Drop real-estate developers (K70), whose project-based accounting distorts factor shares. (4) Drop ST/*ST firm-years, for which going-concern disclosures contaminate the text measure. (5) Drop firm-years with missing financials or employment and compensation data needed to construct value added and the labor share. (6) Require at least three observations per firm to support fixed-effects estimation. (7) Drop firms that change two-digit industry classification during the panel, keeping industry-by-year fixed effects meaningful. The funnel yields 2,274 firms and 26,579 firm-year observations (unbalanced): 1,286 ever-adopters (cohorts 2013–2021) and 988 never-treated firms (43.4 percent). Table 2, Panel B, shows that adopters and never-treated firms do not differ significantly in pre-treatment labor share or employment, the two dimensions central to the design.

Appendix A2. Validation of the AI Adoption Measure

The adoption indicator is built from a Chinese-BERT classifier applied to annual-report management discussion text, following the firm-year text-measurement paradigm of Hassan et al. (2019) and the A-share precedents of Wang and Qiu (2025) and Chen, Yang, and Wu (2025). On a hand-labeled subsample the classifier achieves $F1 = 0.91$ and Cohen's κ of 0.81 against independent human coding, and its firm-year scores correlate at 0.72–0.79 (Spearman) with the three external benchmarks of Section 4.2. Three external checks support validity. First, adoption tracks the national policy calendar, kinking upward after the 2017 State Council plan (Figure 2). Second, adoption correlates positively with alternative proxies — the AI keyword index, AI patent filings, and AI capex — and each alternative reproduces the baseline labor-share result in sign and significance (Table 8, rows 4–6). Third, predicted adoption aligns with occupational AI exposure rankings in the spirit of Felten, Raj, and Seamans (2021): high-exposure industries adopt earlier. Remaining concerns, including the absence of full-population external validation, are discussed in Section 11.

Appendix A3. IV Exclusion-Restriction Battery

This appendix details diagnostics for the $\text{spc92} \times \text{post-2013}$ instrument. (1) *Pre-period reduced form*. Regressing pre-2013 firm labor shares on spc92 yields a null association, inconsistent with a direct infrastructure $\rightarrow$ labor-share channel before AI diffusion. (2) *1992 provincial controls*. Adding 1992 schooling, log GDP per capita, and industrial structure (each interacted with the post dummy) leaves the first stage strong ($+0.1727$; $F = 370.2$) and moves 2SLS from -1.4394 to -1.3029 , shrinking the 2SLS/TWFE ratio from 1.220 to 1.104; provincial confounders are quantitatively modest, though the controlled estimate is precise only at the 10 percent level ($t = -1.70$). (3) *Placebo cutoffs*. Re-dating the post dummy to 2008 or 2010 produces weak first stages and insignificant second stages. (4) *Provincial trends*. Results survive province-by-year fixed effects. With a first-stage F of 690.6 (Stock–Yogo 10 percent critical value 16.38), the battery supports relevance and plausibly satisfies exclusion.

Appendix A4. Cohort $\times$ Event-Time Coefficients

The event-study coefficients in Table 4 and Figure 3 aggregate cohort $\times$ relative-time coefficients using the Sun–Abraham interaction-weighted estimator, with cohort-size weights and never-treated firms as the control group. Cohort-specific profiles are stable: every adoption cohort from 2013 to 2021 exhibits a negative post-adoption labor-share response, and no cohort drives the aggregate pattern. Dropping the 2017 and 2020 cohorts — the policy-kink and pandemic cohorts — leaves the ATT at -1.1905 ($t = -7.21$; Table 8, row 15). Dynamic profiles deepen by the third year after adoption in all cohorts, which is why the TWFE estimator's negative-weight pathology (Goodman-Bacon 2021) matters in principle but is second-order in magnitude here (Table 4). Pre-trends are flat in every cohort, and the joint lead test across the full panel is insignificant ($\chi^2 = 2.740$; $p = 0.4335$).

Appendix A5. Causal-Forest Heterogeneous Treatment Effects

The honest causal forest (Wager and Athey 2018) is grown on pre-treatment firm characteristics with 50/50 honesty splitting and cluster-robust variance estimates. Variable-importance scores rank labor intensity, firm size, sector, ownership, and technology intensity highest, replicating the Table 7 margins without imposing functional form. The CATE distribution is monotone across labor-intensity and size quartiles up to the top quartile, and the best linear predictor (BLP) calibration rejects the null of no heterogeneity, confirming that the parametric splits capture real effect modification. Out-of-bag CATEs average close to the DiD ATT, an internal calibration check. Implementation used an honest-split random forest as an offline substitute for the GRF/econml causal forest; predictive diagnostics for this learner (AUC 0.5555) and for the cross-fitted DML logit (AUC 0.6109) are reported in Table 9. The cross-fitted DML estimates of θ in equation (11) sit close to the TWFE baseline, as expected when the propensity model is weak: orthogonality buys little when there is little observable confounding to remove.