NBER WORKING PAPER SERIES

AI AS SHARED BUSINESS INFRASTRUCTURE: FIRM GROWTH AND THE DISTRIBUTION OF GAINS

Panle Jia Barwick
Lu Fang
Jinhan Liu
Huilan Xu
Zhe Yuan

Working Paper 36003
http://www.nber.org/papers/w36003

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2026

The paper only reflects the opinion of the authors and all remaining errors are the responsibility of the authors. We are grateful to various seminar participants for valuable comments and discussions. Lu Fang acknowledges financial support from the National Natural Science Foundation of China [Grants 72501258 and 72192803], the Major Special Project in Philosophy and Social Sciences of the Ministry of Education of China [Grant 2026JZDZ020]. Zhe Yuan acknowledges financial support from the National Key Research and Development Project [Grant 2024YFB3312900], the National Natural Science Foundation of China [Grants 72673168 and 72141305], Major Project of the National Social Science Fund of China [Grant 25 & ZD149 and 26ZDA030]. We thank various conference and seminar participants for their helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Panle Jia Barwick, Lu Fang, Jinhan Liu, Huilan Xu, and Zhe Yuan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

AI as Shared Business Infrastructure: Firm Growth and the Distribution of Gains
Panle Jia Barwick, Lu Fang, Jinhan Liu, Huilan Xu, and Zhe Yuan
NBER Working Paper No. 36003
November 2026
JEL No. D22, L25, L81, L86, O33

ABSTRACT

Generative AI improves performance on individual tasks, but whether those gains translate into firm growth, and for which firms, remains an open question. We study a leading Chinese e-commerce platform that in 2024 embedded a generative AI system in the operating software firms use to run their businesses, drawing on centrally developed models and platform-wide data to deliver ready-to-use applications to every firm. Complete records of each firm's AI use, linked to its monthly operations and sales, let us trace adoption through changes in business capabilities, such as reaching customers and introducing new products, to sales growth. Using a matched difference-in-differences design and focusing on firms with annual sales below RMB 5 million, we find that adoption raises monthly sales by 51.6%, primarily through a 47.0% increase in the number of orders; instrumental-variable estimates confirm these results. Adopters reach more customers, update their product lines faster, and deliver better product and service quality, with the largest proportional gains for the smallest firms. Two mechanisms help explain the results. First, AI lowers the cost of routine work: adopters produce more product content and change listings more often while spending less staff time per order, and the gains are largest for firms that faced the highest costs relative to sales before adoption. Second, AI reduces information frictions: adopters post richer product information, and the gains are largest in markets where demand is hard to predict and matching quality is poor, especially for small firms. At the market level, product categories with greater pre-existing information frictions see more sales by existing firms, more entry, and more new products. AI expands markets rather than merely shifting sales from non-adopters to adopters.

Panle Jia Barwick
University of Wisconsin - Madison
Department of Economics
and CEPR
and also NBER
panle.barwick@gmail.com

Lu Fang
Zhejiang University of Finance
and Economics
School of Management
fl_fanglu@163.com

Jinhan Liu
Zhejiang University
School of Economics
ljhcy0616@gmail.com

Huilan Xu
Zhejiang University
yyxuhuilan@gmail.com

Zhe Yuan
Zhejiang University
School of Economics
yyyuanzhe@zju.edu.cn

1 Introduction

Generative AI can substantially improve performance on individual tasks (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Cui et al., 2026). Whether these gains translate into firm growth is an open question. To turn task-level gains into firm growth, firms must build AI into recurring activities in ways that strengthen their business capabilities: reaching customers, introducing new products, and maintaining quality. The required investments in models, data, workflows, and organization have so far favored large firms (Brynjolfsson, Rock, and Syverson, 2021; McElheran et al., 2024; Babina et al., 2024; Acemoglu et al., 2025). AI is built at scale, and to date its benefits have flowed mostly to firms with scale.

Yet a growing share of AI tools reaches firms not through their own investments but through external providers: cloud services, enterprise assistants such as ChatGPT Enterprise, and platform-integrated assistants such as Shopify Sidekick (Hagiu and Wright, 2025; Brynjolfsson, Li, and Raymond, 2025; Demirer, Fradkin, and Tadelis, 2026). When a provider absorbs the scale-intensive investments and delivers ready-to-use applications, task-level gains may translate into firm-level gains for firms that could never have made those investments themselves. Whether, and for whom, AI translates into growth may therefore depend not only on what the technology can do but on how it is supplied. We study the most integrated form of this arrangement, which we call *AI as shared business infrastructure*: a leading Chinese e-commerce platform that develops AI models centrally, combines platform-wide information with each firm's own data, embeds ready-to-use applications in the software firms already use to run their stores, and makes them available to every firm on the platform.

Testing whether, and for whom, task-level gains translate into firm growth requires observing three things together, and across the full distribution of firm size: which firms actually use AI and for what purpose, how their operating activities and business capabilities change, and what happens to their performance. Existing studies, however, often infer adoption from surveys, AI-related job postings, patents, or exposure measures (Damioli, Van Roy, and Vertesy, 2021; Czarnitzki, Fernández, and Rammer, 2023; da Silva Marioni, Rincon-Aznar, and Venturini, 2024; Babina et al., 2024;

McElheran et al., 2025; Eisfeldt et al., 2026), proxies that capture large firms well and small firms poorly: a firm with five employees that uses AI to write its listings posts no AI job advertisement and files no patent. The platform's records let us see every firm's AI requests, the functions it used, and its monthly operations and transactions. We therefore observe the entire chain, from use to capabilities to growth, including for small firms that existing measures often miss.

Which firms gain most is not obvious. Shared infrastructure removes the fixed costs and integration burdens that weigh most heavily on small firms (DeStefano, Kneller, and Timmis, 2025). But firms with more knowledge, better managers, and stronger organizations may still extract more from the same tool (Cohen and Levinthal, 1990; Bresnahan, Brynjolfsson, and Hitt, 2002), and Otis et al. (2026) find that an AI business assistant helped high-performing entrepreneurs and hurt low-performing ones, because the two groups chose to implement different parts of otherwise similar advice. Firm-level gains also need not imply market expansion: adopters may grow by creating new activity, by taking business from non-adopters, or both. Assessing the consequences of shared AI infrastructure therefore requires evidence on three margins: changes in firms' capabilities and performance, how those changes vary across the size distribution, and market-wide adjustment.

To guide the empirical analysis, we develop a simple model in which firms with different initial capabilities adopt AI that lowers the inputs routine tasks require and supplements firms' own information with platform-wide signals. Adoption raises sales through both channels, and by more in markets where pre-existing information frictions are more severe. In the model, however, the direction of the size gradient is ambiguous. Because the platform delivers the same tools to every adopter, AI partly duplicates what high-capability firms already have and fills larger gaps for low-capability firms. But if the value a firm extracts from AI rises steeply enough with its existing capabilities, proportional gains instead rise with firm size. This happens, for example, when better-managed firms are better at selecting and implementing AI output. Which force dominates is therefore an empirical question.

The platform we study began testing the AI system in March 2024. The system has two broad functionalities: content generation and business operation. The content-generation functionality,

launched in April 2024, generates and edits product images, videos, titles, and other listing materials. The business-operation functionality, launched in July 2024, helps firms follow platform rules, produces daily performance reports and competitor benchmarks, suggests products and suppliers, supports promotions, and assists with customer service and fulfillment. Together, the two functionalities span six broad business functions. Each firm receives one month of free access and then pays RMB 99 (about $14) per month.

The platform's records are linked at the firm-month level and span October 2023 to July 2025. The usage logs record who adopted, when, and which functions they used; the operating data cover each firm's sales, customer traffic, product listings, quality measures, staffing, and firm characteristics. We study firms that sell physical goods and had pre-launch annual sales of at most RMB 5 million. Our baseline sample further restricts attention to continuing firms, those with positive sales in every month; as a robustness check, we also include firms with zero sales. For each month, we match firms that adopted both tools in that month to similar firms that never adopted during the sample period, based on pre-adoption characteristics. Because a non-adopter can serve as a control for more than one cohort, the matched baseline sample contains 148,770 adopters and 77,347 distinct non-adopters. Pooling the monthly cohorts yields more than 6 million firm-cohort-month observations.

We estimate a matched-cohort difference-in-differences design with firm-by-cohort and cohort-by-calendar-month fixed effects. In other words, the estimates compare changes within a firm in the same cohort to changes for its matched controls over the same calendar months. We use a balanced event study to examine treatment dynamics and assess differential pre-adoption trends between adopters and their matched controls. Because firms choose whether and when to adopt, we also report instrumental-variable estimates to address potential endogeneity of the adoption decision. The instrument predicts a firm's adoption from three conditions fixed before the AI existed or determined outside the firm's own market: how much the firm relied on the tasks the AI now performs, how digitally developed its province was, and how quickly AI was spreading in neighboring product categories. We also consider alternative estimators for staggered adoption.

Our matched-cohort estimates imply substantial improvements in firm performance. Monthly sales rise by 51.6% and orders by 47.0%, with little change in average order price: adopters grow by completing more transactions, not by charging more per transaction. Instrumental-variable estimates yield the same qualitative conclusions. The gains exhibit a pronounced size gradient: monthly sales rise by 61.5% among firms in the bottom quartile of pre-launch sales, compared with 10.3% in the top quartile.

The sales gains rest on broader improvements in how firms operate. We examine three business capabilities: reaching customers, updating product offerings, and maintaining quality. Adopters reach more customers: page views and new visitors rise by roughly 45%, and sales to new customers rise by 75.6%. Adopters adapt their product lines faster, introducing 66.2% more new products and earning 50.5% more sales from newly introduced products. Importantly, this growth does not come at the expense of quality: rates of firm-related refunds, negative reviews, and firm-liability disputes all fall, while the 24-hour pickup rate rises and delivery time falls. Here, too, the improvements are more pronounced among smaller firms.

Evidence points to two mechanisms. First, platform AI lowers the cost of routine work. Adopters carry out more routine work with higher labor efficiency: the number of product images rises by 62.6% and listing updates per product by 124.8%, while staff time per order falls by 5.6%. Before the AI, firms could buy comparable services from external providers, but the fees are large relative to a small firm's revenues. At market prices, producing the same content would cost the median firm in the bottom size quartile 22.7% of its annual sales, compared with 0.4% in the top quartile, and outsourcing operating support would cost more than most small firms earn in a year. For these firms, the AI thus provides services that were previously out of reach. Consistent with this channel, sales gains are largest for firms whose operating costs were highest relative to their sales before adoption.

Second, platform AI reduces information frictions. The business operation tool provides each firm with ready-to-use analyses of demand, competitors, and its own operations, analyses that large firms produce with dedicated staff and small firms often do without. Consistent with this channel,

adopters post richer product information: titles lengthen by 20.2%, and the share of product attributes reported rises by about 30%. Moreover, sales gains are largest where pre-AI frictions were most severe, whether measured by refund rates, which signal poor matches between products and buyers, or by the volatility of firms' sales, and the relationship is strongest among smaller firms.

A natural concern is that adopters gain by taking sales from non-adopters. Because the tools were available to every firm on the platform, no product category can serve as an untreated control. As AI tools matter most where information frictions were most severe, we compare high- and low-friction categories. The two moved in parallel until the platform launched the AI tools, and then they diverged: high-friction categories saw larger increases in total sales, incumbent sales, firm entry, entrant sales, and new-product introductions. If adopters were simply taking sales from other firms, category sales would not grow faster where the tools matter most, and new firms would not be drawn in. Additional tests based on historical counterfactuals, firm exit, and competitive exposure also provide little evidence that adopters' gains come at the expense of non-adopters. The evidence points instead to market expansion.

Literature Review and Contribution This paper contributes to three related literatures. First, we provide direct evidence on whether task-level gains from AI translate into firm growth. Experiments establish large task-level productivity effects (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Cui et al., 2026), but firm-level evidence has been held back by measurement. Existing measures capture different dimensions of AI activity rather than realized use (Babina, 2026): surveys on self-reported adoption (Czarnitzki, Fernández, and Rammer, 2023; McElheran et al., 2025), AI-related hiring (Alekseeva et al., 2021; Babina et al., 2024), and patents (Damioli, Van Roy, and Vertesy, 2021; da Silva Marioni, Rincon-Aznar, and Venturini, 2024). These measures also underrepresent small firms, which rarely hire AI specialists, patent, or go public. Prior work argues that translation requires firms to integrate AI into recurring activities and adapt their workflows (Brynjolfsson, Rock, and Syverson, 2021; McElheran et al., 2025; Krakowski et al., 2026). We observe that integration directly, in addition to AI adoption and usage, for every firm

on the platform, and show that it changes routine operations, strengthens customer reach, product adaptation, and quality, and raises orders and sales.

Second, we show that economies of scale in AI production need not translate into scale-biased downstream gains when AI is supplied as integrated platform infrastructure. Building frontier AI involves large economies of scale (Varian, 2019; Goldfarb and Trefler, 2019; Korinek and Vipra, 2025), and when firms make the complementary investments themselves, both adoption (McElheran et al., 2024; Acemoglu et al., 2025) and measured gains (Bäck et al., 2022; Babina et al., 2024) concentrate among larger firms. External provision can lower these barriers: access to cloud computing and generative AI has been linked to more entry and faster growth among young and small firms (Lu, Phillips, and Yang, 2024; Cai et al., 2025; DeStefano, Kneller, and Timmis, 2025). We study provision in its most complete form, in which the platform absorbs not only development costs but much of the adaptation work that normally falls on adopters, and find the largest proportional gains among the smallest firms. Scale economies in producing AI need not become scale advantages in using it.

Third, we contribute to research on whether AI-driven gains create new activity or reallocate it from competitors (Bloom, Schankerman, and Van Reenen, 2013; Horn, Sacarny, and Zhou, 2022; Lu, Phillips, and Yang, 2024; Athey and Scott Morton, 2026). Market-level evidence so far comes from creative-goods markets, where AI-generated content expands output but crowds out human producers (Goldberg and Lam, 2025). Because we observe adoption and outcomes across an entire retail platform, we can study firm- and market-level responses together. Categories where AI should matter most grow faster in total sales, entry, and new products, a pattern consistent with market expansion rather than pure reallocation.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and presents the conceptual framework. Section 3 describes the data and empirical strategy. Section 4 reports the effects on firm performance and business capabilities, checks their robustness, examines potential spillovers, and presents the market-level evidence. Section 5 examines heterogeneity by initial firm size. Section 6 explores mechanisms. Section 7 concludes.

2 Institutional Setting and Conceptual Framework

This section describes the AI system studied in this paper and develops a conceptual framework. The framework yields four testable predictions about firm growth, the two channels through which AI operates, and market expansion. It also shows why the direction of the size gradient cannot be determined ex ante.

2.1 Platform AI System

The AI system we study was introduced by one of the world's largest digital platforms.[1] It connects millions of firms with hundreds of millions of consumers across a broad range of product categories. The AI system combines two broad functionalities: content generation and business operation. Both are integrated into firms' operating workflows and draw on firm-specific and platform-wide information. Together, they encompass six functions that automate or assist with routine business tasks, from content production and data analysis to customer service and fulfillment.

The *Content-Generation functionality* supports two broad business functions. *Content Editing* revises existing listing materials, such as images, videos, titles, layouts, and promotional text, while *Content Generation* creates new listing content. Firms can use these functions for individual tasks or batch production. By automating parts of content production, these tools reduce the production time and specialized inputs firms would otherwise need to provide internally or buy from outside agencies.[2] Appendix Figure A.1 illustrates representative examples of the content generation tool.

The *Business-Operation functionality* supports four broad business functions. *Operational Guidance* helps firms follow platform-specific rules and address operational issues. Unlike a traditional help center that provides static, general information, it provides contextualized answers

[1]The platform is anonymized under a data-use agreement.

[2]Conventional product photography can be costly and time-consuming, often requiring human models, photo shoots, and subsequent editing over several days, whereas the AI tool can generate listing-ready images from a few input photos in roughly 20 seconds. A similar contrast applies to video production: the tool can generate promotional videos from product images within three to five minutes. Appendix F provides the reference prices used to construct cost-equivalent benchmarks for these activities.

and recommended actions based on firm-specific questions and operating conditions. *Data Analytics* combines firm operating data with platform-wide market information to produce performance reports, competitor benchmarks, market trend analysis and practical recommendations. *Market & Product Management* assists with product selection, sourcing, promotion, and customer outreach. *Service & Fulfillment* supports customer interactions, fulfillment and after-sales service. Once configured, the system can automatically perform recurring tasks such as customer messaging, logistics tracking, and daily performance reporting. By automating routine operational tasks and turning firm and platform data into tailored diagnoses and recommendations, these tools reduce the information-processing and managerial effort and support better-informed decisions. Appendix Figure A.2 provides representative examples of the business operation tool.[3]

AI Rollout and Adoption The platform began piloting the AI system in March 2024, introduced the content-generation functionality in April, and added the business-operation functionality in July. Each firm receives one month of free access, followed by a monthly fee of RMB 99 (about $14).

Figure I shows rapid diffusion of the AI system. By July 2025, more than half of the platform firms observed during the study period had used at least one of its two functionalities. The platform rollout of the system is not staggered. However, firms choose whether and when to adopt, resulting in staggered adoption timing in our firm-level analysis. Appendix Figure A.3 illustrates AI use by functionalities and initial firm-size quartile. *Operational Guidance*, *Content Generation*, and *Content Editing* account for the largest shares of firm-initiated requests.[4]

[3]For example, *Operational Guidance* covers platform fees, promotional-event eligibility, listing and quality rules, and violation procedures, and flags emerging issues such as shipment delays. *Service & Fulfillment* handles customer inquiries, order and logistics tracking, returns, and repairs.

[4]Among users, approximately 20% used only content generation, 30% used only business operations, and slightly more than half used both. Smaller firms devote a larger share of their requests to *Operational Guidance*, *Content Editing*, and *Market & Product Management*, whereas larger firms allocate a larger share to *Content Generation*, *Data Analytics*, and *Service & Fulfillment*.

2.2 Conceptual Framework and Predictions

The institutional features described above suggest two channels through which platform AI can improve firm performance and capabilities. First, it automates and supports routine business activities by reducing the labor, time, and expertise they require (Acemoglu and Restrepo, 2018, 2019). Second, it improves decision-making by combining firm data with platform-wide information and turning these data into useful operational insights and guidance (Garicano, 2000; Milgrom and Tadelis, 2019; Jovanovic and Rousseau, 2026).

We capture these channels in a static model of monopolistic competition with heterogeneous firms.[5] Firms differ in their pre-AI business capability k_i, including operating knowledge, proprietary information, managerial capacity, and task efficiency. Higher-capability firms have better information, require fewer task inputs, and have higher baseline sales. This monotonic relationship motivates our use of pre-launch annual sales to rank firms by initial size, while recognizing that sales are only a noisy proxy for the latent capability index.

In the model, AI reduces the inputs that routine tasks require and improves the quality of firms' decisions and listings, partly directly and partly by supplying information that firms would not otherwise have. The key question is how these task-level improvements translate into firm-level gains when firms start from different capabilities. Two forces pull in opposite directions. The first is duplication: because high-capability firms already have better information and lower task inputs, standardized platform tools partly duplicate what they have and fill larger gaps for lower-capability firms. If every adopter realized the same increment from these tools, duplication would be the only force, and proportional gains in capability, sales, and operating-cost savings would be largest for initially lower-capability firms. The second is complementarity: what the platform delivers is standardized, but what a firm realizes from it need not be. Firms with stronger managers and organizations may be better at choosing which AI outputs to use, adapting them to their products, and acting on AI recommendations (Otis et al., 2026). If complementarity is strong enough, pro-

[5] Each product category contains a continuum of firms that supply differentiated varieties under CES demand. Firms set prices to maximize profits.

portional gains instead rise with initial capability (Appendix B, Lemma 3). The model does not determine which force dominates, so we treat the direction of the size gradient as an empirical question (Section 5). Appendix B presents the full model, comparative statics, and proofs.

The model yields four predictions that organize the empirical analysis. The first three hold category-level aggregates fixed; the fourth lets them adjust.

Prediction 1. *(Firm growth)* Platform AI improves firms' sales performance and business capabilities.

Prediction 2. *(Operating-cost channel)* Platform AI lowers the cost of routine work, allowing firms to perform more business activities and increase sales.

Prediction 3. *(Information-friction channel)* Platform AI reduces information frictions by supplementing firms' private information with platform-wide signals, improving decision-making, information provision, and sales. These gains are proportionally larger in markets with more severe pre-existing information frictions.

When many firms adopt, lower operating costs and improved capabilities reduce the effective price of the category composite. Because platform information is most valuable where information frictions are severe, those markets should expand most.

Prediction 4. *(Market expansion)* Broad adoption of platform AI increases category quantity, with larger gains in categories with more severe pre-existing information frictions. When category demand is sufficiently elastic, category sales also increase.[6]

3 Data and Empirical Strategy

We measure AI adoption directly from usage records and estimate its effects at three levels: the firm, its competitors, and the market. This section describes the data and sample, the matching

[6]The number of orders is not a clean measure of category quantity because orders may contain multiple units and products differ within categories. We test Prediction 4 using category sales, incumbent sales, firm entry, entrant sales, and new-product introductions.

procedure, and the empirical strategies.

3.1 Data Description

We assemble a firm-month panel from October 2023 to July 2025 using four sources of proprietary platform data. First, AI-use records identify whether and when firms adopted AI and which functions they used. These records capture firm-initiated content-generation tasks and business-operation requests, excluding downstream outputs and backend model or API activity. Second, monthly transaction records report sales, order volume, and average order price. Third, monthly operating records cover customer reach, product offerings, and product and service quality. Fourth, firm characteristics include tenure, geographic location, and product category. Together, these data form the basis of our firm-level analysis.

We construct a category-month panel for the market-level analysis. The platform classifies products in a four-level hierarchy, and we aggregate the underlying records to the category-month level at the most detailed level, Level 4.[7] We also use pre-sample platform records to construct baseline measures, the instruments, and province-level digitization measures from 2023.

Sample Restrictions For the firm-level analysis, we first draw a random 80% sample of all firms on the platform and then impose three sample restrictions before conducting propensity score matching. First, we require firms to record at least one transaction in every sample month, so the baseline firm-level analysis focuses on intensive-margin effects among continuing firms rather than changes driven by entry and exit. Appendix D.5.1 incorporates firms with zero-sales in some months, so that adoption can also affect whether a firm makes positive sales. We examine firm entry and exit separately in the market-level analysis.

Second, we exclude firms with more than RMB 5 million in sales over the 12 months before

[7]Level 1 comprises approximately 24 broad product sectors, Level 2 more than 100 categories, Level 3 more than 2,000 subcategories, and Level 4 more than 20,000 detailed product categories. For example, the hierarchy runs from "clothing" (Level 1) to "women's apparel" (Level 2), "skirts" (Level 3), and "midi skirts" (Level 4).

the April 2024 launch.[8] These firms are a thin upper tail of the size distribution, accounting for about 1% of all firms on the platform. They adopt the platform AI at much higher rates, leaving few comparable non-adopters to serve as controls. Many also rely on other AI tools and in-house technical staff, so adoption of the platform tool is a poor measure of their AI use. Our baseline estimates therefore apply to firms below this threshold.

Third, we keep only firms selling physical goods, excluding virtual goods, pharmaceuticals, and services, whose production, transaction, and fulfillment processes differ substantially from those of conventional retail. As a share of all firms on the platform, the sample shrinks from 80% (the random draw) to 18.4% after requiring positive sales in every month, to 16.6% after excluding firms above RMB 5 million, and to 11.3% after restricting to physical goods. The activity requirement removes the most firms because many firms on the platform sell only intermittently; Appendix D.5.1 examines the effect of AI on these firms. The baseline sample thus consists of continuing physical-goods firms with pre-launch annual sales of at most RMB 5 million.

3.2 Treatment Definition and Matching

We define treatment from firms' recorded use of the two AI functionalities. A firm adopts a functionality in the first month in which it initiates a request using it. Because the system we study bundles the two functionalities, treatment requires both: *treated firms* adopt Content Generation and Business Operation no more than one calendar month apart, and the earlier of the two months is the treatment month.[9] *Control firms* use neither functionality at any point during the sample period. Firms that adopt only one functionality, or adopt both more than one month apart, are excluded from the baseline analysis; Appendix D.5.2 examines single-functionality adopters separately. This definition lets us estimate the effect of adopting the integrated system rather than either

[8]The threshold follows China's VAT classification: retail enterprises with annual sales of RMB 5 million or less are small-scale VAT taxpayers, and those above it are general VAT taxpayers.

[9]Because some treated firms adopt the second functionality one month after the first, we re-estimate the baseline specification excluding the treatment month for both treated firms and their matched controls. The post-treatment period then begins only after the treated firms have adopted both functionalities. The estimates are similar to the baseline results.

functionality alone, and the one-month window keeps the timing of adoption comparable across treated firms. Joint adoption is also the common pattern: more than half of AI users adopted both functionalities during the sample period, and the share has continued to rise since.

We group treated firms into 12 monthly *adoption cohorts* based on their treatment month, from July 2024 through June 2025. Each cohort consists of the treated firms sharing the same treatment month. Cohort membership reflects when a firm actually adopted, not when the tools became available.

We then construct a matched control group for each cohort. We estimate propensity scores from pre-adoption characteristics (tenure, primary product category, firm flagship-store type, sales, orders, and the number of active products) on the sample of the cohort's treated firms and the full pool of never-adopters, and we match each treated firm, with replacement, to the never-adopter with the closest score.[10] Each treated firm is thus paired with one control. Because matching is with replacement, one never-adopter can serve as the control for multiple treated firms. Repeated matches to the same control within a cohort are consolidated and accounted for using matching-frequency weights in the regressions. Following Cengiz et al. (2019), we stack the 12 cohort-specific matched samples into a single dataset, with each adopter–control pair observed over the common October 2023–July 2025 window. The matched sample contains 148,770 adopter–control pairs and 77,347 unique never-adopters.

Summary Statistics Table I reports summary statistics for the matched sample; Appendix Table A.1 defines the variables. After matching, treated firms and their controls are similar on pre-adoption characteristics: all absolute standardized mean differences in Appendix Table A.2 are below 0.10, the conventional threshold for balance. As discussed in Section 3.1, the sample consists mostly of small firms, though with wide variation in size. Pre-launch annual sales, measured over the 12 months before April 2024, have 25th, 50th, and 75th percentiles of RMB 24,460, RMB

[10]Primary product category is the unique Level-1 category generating the largest share of a firm's sales over the preceding year. Flagship-Store Type is an indicator equal to one for flagship stores and zero for independent third-party firms. Time-varying covariates are averaged over the cohort's pre-treatment period.

63,088, and RMB 147,716, and these quartile cutoffs define the four size groups used in the heterogeneity analysis.

3.3 Empirical Strategy for Firm-Level Analysis

Our firm-level analysis uses three complementary research strategies. A matched difference-in-differences design estimates the average effect of adoption; a balanced event study traces its dynamics and tests for differential pre-trends; and an instrumental-variable strategy addresses the concern that firms' adoption decisions could be endogenous.

Matched Difference-in-Differences Design Let $Treat_{ig}$ equal one if firm i is treated in cohort g and zero if it is a never-adopter matched to that cohort, and let $Post_{gt} = \mathbf{1}\{t \geq G_g\}$, where t indexes calendar month and G_g denotes the adoption month of cohort g. We estimate

$$y_{igt} = \beta \left(Treat_{ig} \times Post_{gt} \right) + \mu_{ig} + \tau_{gt} + \varepsilon_{igt}, \tag{1}$$

where y_{igt} is the outcome of firm i in month t within cohort g. Continuous nonnegative outcomes are transformed using $\log(1 + x)$ whereas rate and index outcomes remain in levels unless otherwise noted. For five outcomes with a substantial mass at zero (new products, new categories, new-product sales, images, and modifications per product), we report Poisson pseudo-maximum likelihood (PPML) estimates and their implied percentage effects in the main text, with corresponding $\log(1 + x)$ OLS-DID and IV-2SLS estimates in Section 4.3 for robustness. Standard errors are clustered at the Level-2 product-category level, which the platform commonly uses to classify firms by their primary business activity. This allows for correlated shocks across firms operating in the same product market.

The firm-by-cohort fixed effects, μ_{ig}, absorb time-invariant differences between firms within a matched cohort. The cohort-by-calendar-month fixed effects, τ_{gt}, give each cohort stack an unrestricted calendar-time path, absorbing shocks common to its firms in a given month. Identification

thus comes from comparing changes for adopters with changes for their matched never-adopters within the same cohort and over the same calendar months. Under no anticipation, no cross-firm spillover, and conditional parallel trends within matched cohorts, β is a weighted average of cohort-specific difference-in-differences estimates for adopters. We assess possible violations of no spillovers in Section 3.4. The baseline DID design uses all twelve cohorts that adopt between July 2024 and June 2025. Because the panel ends in July 2025, later cohorts contribute fewer post-treatment months, so β averages effects across post-adoption horizons rather than measuring the effect at a common event time.

Event Study To trace dynamics and test for differential pre-trends, we estimate a balanced event study on the seven cohorts that adopt between July 2024 and January 2025, using a common six-month window before and after adoption so that every event-time coefficient is estimated from the same set of cohorts rather than a changing cohort composition (Wing, Freedman, and Hollingsworth, 2024):

$$y_{igt} = \sum_{\substack{\ell=-6 \\ \ell\neq-1}}^{6} \beta_\ell \left[Treat_{ig} \times \mathbf{1}\{t - G_g = \ell\} \right] + \mu_{ig} + \tau_{gt} + \varepsilon_{igt}. \tag{2}$$

The specification retains the same fixed effects as in equation (1). For matched controls, event time runs relative to the treatment month of their cohort. We omit $\ell = -1$, so that β_ℓ measures the treated-control difference at event time ℓ relative to the month before treatment.

Instrumental-Variable Strategy Firms choose whether and when to adopt AI, so adoption may respond to time-varying shocks that also affect performance: selection is positive if firms adopt in anticipation of growth, negative if operating difficulties prompt adoption. Matching, fixed effects, and the event-study evidence mitigate this concern but cannot rule it out. We therefore complement the matched difference-in-differences design with an instrumental-variable strategy.

Following differential-exposure and shift-share designs (Duflo, 2001; Borusyak, Hull, and Jaravel, 2022; Barwick et al., 2026), we construct an instrument that interacts three components fixed

before the AI launch or determined outside the firm's own product market:

$$Z_{it} = D_i \times R_{p(i)} \times A_{-c(i),h(i),t}, \tag{3}$$

where D_i is firm i's pre-launch reliance on tasks the AI now performs, measured by product-listing modifications and human customer-service time per order. $R_{p(i)}$ is pre-launch digital readiness in firm i's province, combining computer use, website adoption, and e-commerce participation. $A_{-c(i),h(i),t}$ is the month-t AI adoption rate in the other Level-2 categories under firm i's Level-1 parent category; it excludes the firm's own category and thus captures diffusion in related rather than own markets, similar in spirit to the related-industry exposure measure of Helm (2020). As AI spreads through related markets, firms that rely more on AI-relevant tasks and operate in more digitally ready provinces should be more likely to adopt.

We estimate the following first-stage equation:

$$AI_{it} = \lambda Z_{it} + \mu_i + \tau_t + \nu_{it}, \tag{4}$$

where AI_{it} equals one during and after firm i's treatment month, and μ_i and τ_t are firm and calendar-month fixed effects.[11] The second stage is

$$y_{it} = \beta \widehat{AI}_{it} + \mu_i + \tau_t + \varepsilon_{it}. \tag{5}$$

The exclusion restriction requires that related-market AI diffusion affect a firm's performance only through its own adoption: related-market adoption may encourage the firm to adopt, but may not affect its outcomes directly. The main threats are competitive spillovers and correlated regional or product-market shocks. The leave-out construction prevents the firm's own adoption from mechanically entering the instrument and excludes adoption by competitors in its own Level-2 category.

[11]The IV uses the same firms as the matched DID (148,770 adopters and 77,347 never-adopters) but is estimated on the simple firm-month panel: each firm appears once per month, with no cohort stacking, so the equations carry no cohort index g.

Industry-specific AI training facilitates diffusion across related product markets, while routine business support is generally not allocated by industry. We assess sensitivity to common regional and product-market shocks using province-by-month and Level-2 category-by-month fixed effects (Appendix Table C.1); Section 3.4 tests for spillover effects directly. Following Nunn and Qian (2014), we further interact firms' pre-AI sales, product assortment, and tenure with calendar-month indicators to allow for differential growth paths associated with predetermined firm characteristics. The IV estimates remain robust, mitigating concerns that the instrument captures pre-existing differences in firm growth rather than AI adoption.

The IV identifies effects for firms whose adoption responds to the instrument and need not match the weighted average identified by the matched DID; we report the two sets of estimates side by side in the main tables. As before, standard errors are clustered at the Level-2 category level. Appendix C details the construction of the instrument.

Sections 4.1 and 4.2 report the matched-DID, event-study, and IV estimates, and Section 4.3 reports robustness checks, including the Borusyak, Jaravel, and Spiess (2024) estimator for staggered adoption and synthetic difference-in-differences.

3.4 Spillover Tests and Market-Level Design

Two questions remain after the firm-level analysis: whether competition from adopters contaminates the treated–control comparison, and whether AI expands markets or merely reallocates them. We address the first with three spillover tests and the second with a category-level event study.

Competitive Spillover Tests The firm-level DID requires that one firm's adoption not affect other firms' outcomes. If adopters take demand from non-adopters in the same market, matched controls perform worse because of AI, violating the stable unit treatment value assumption (SUTVA) (Athey, Eckles, and Imbens, 2018) and leading the treated–control gap to overstate the effect of a firm's own adoption. We conduct three tests to evaluate this.

First, we construct a historical counterfactual for the July–December 2024 comparison sample.

For each control firm, we replace its post-period outcomes with outcomes from the same months of 2023, adjusted for aggregate growth, and re-estimate the baseline design on the same sample. Because these historical outcomes predate the AI system, they are immune to AI-driven competition: if such competition depresses contemporaneous control outcomes, the baseline estimate should exceed the historical-counterfactual estimate.

Second, we test whether greater category-level AI diffusion raises firm exit, the most severe form of competitive displacement. We examine exit overall and among non-adopters, since aggregate rates may mask offsetting effects between the two groups. We estimate

$$ExitRate_{ct} = \beta AdoptionRate_{ct} + \alpha_c + \delta_{g(c),t} + \varepsilon_{ct}, \tag{6}$$

where c indexes detailed product categories, t indexes months, and $g(c)$ denotes the higher-level category containing c. A firm is classified as exiting category c in month t if it has positive sales in that category in month $t-1$ but no sales in months t, $t+1$, or $t+2$. $ExitRate_{ct}$ is the share of firms active in category c in month $t-1$ that satisfy this exit definition. The non-adopter exit rate is defined analogously, restricting the risk set to firms that had not adopted AI by month t. $AdoptionRate_{ct}$ is the cumulative share of active firms in category c in month t that have adopted AI. Category fixed effects, α_c, absorb time-invariant differences across categories, while higher-level-category-by-month fixed effects, $\delta_{g(c),t}$, compare detailed categories within the same broader market and month.

Third, we test whether the treated-control gap widens with exposure to AI-enabled competitors, as business-stealing would predict. We interact $Treat \times Post$ with two measures of Level-2 category-level AI exposure: the cumulative number of adopters, among large firms and among all firms, divided by the current number of active firms in the category.[12]

All three exercises target negative business-stealing spillovers; they do not rule out positive spillovers across firms. If, for instance, adopters' improved products and content draw additional

[12]Large firms have annual sales above RMB 5 million. All regressions include the constituent lower-order terms of the triple interaction; their coefficients are omitted for brevity.

consumers to the platform, non-adopters could benefit as well (Section 4.4). Such spillovers would raise control firms' outcomes, leading the matched comparison to understate the effect of adoption; in that case, our firm-level estimates are conservative.

Market-Level Design The spillover tests ask whether competition contaminates the firm-level comparison; they do not establish whether AI expands market activity or merely reallocates it. To examine expansion, we use the Level-4 category-month panel. The platform's extensive transaction records allow us to measure market activity at this granular level, capturing changes across narrowly defined product categories. Because the AI system was available platform-wide, no category went untreated, so we exploit cross-category differences in treatment intensity with a continuous-exposure event study (Callaway, Goodman-Bacon, and Sant'Anna, 2026). We proxy baseline information frictions with 2023 refund rates: higher refund rates indicate poorer product-consumer matching and thus greater scope for AI to improve it.

Let BRR_c denote the 2023 refund rate in Level-4 category c, measured as refunded orders divided by total orders. We estimate:

$$y_{ct} = \sum_{\substack{k=-5 \\ k\neq -1}}^{16} \beta_k \left[BRR_c \times \mathbf{1}\{t - T_0 = k\}\right] + \mu_c + \tau_t + \varepsilon_{ct}, \tag{7}$$

where T_0 is March 2024, y_{ct} is a category-level outcome (e.g., total sales, incumbent sales, new-firm entry, entrant sales, new-product introductions, and market concentration), μ_c denotes category fixed effects, and τ_t denotes calendar-month fixed effects. We omit February 2024 ($k = -1$). All outcomes except HHI are transformed using $\log(1 + x)$; HHI is measured in levels. Standard errors are clustered by category.

The coefficients β_k capture differential changes across categories with different baseline exposure. A causal interpretation requires parallel trends absent AI diffusion and no other post-March shocks systematically related to baseline refund rates. Because BRR_c measures scope for AI gains rather than realized adoption, the design identifies differential post-pilot changes in higher-friction

categories, not an effect of category-level adoption or a platform-wide average.

4 Effects of AI Adoption on Firm Performance

This section establishes the paper's first set of results: AI adoption improves firm performance and the business capabilities that support it (Prediction 1). We then show that these results are robust to alternative estimators, samples, and outcome definitions, and that the gains reflect market expansion rather than reallocation across firms (Prediction 4).

4.1 Sales Performance

AI adoption substantially improves firm performance. Figure II shows the dynamic effects from the balanced event study. Sales and orders rise upon adoption and continue to grow thereafter, with no differential pre-trends; average order price changes little throughout. Panel A of Table II quantifies the average effects from the baseline matched-DID estimates: monthly sales increase by 51.6% and orders by 47.0%, with no statistically significant change in average order price.[13] Firms thus grow by completing more transactions, not by charging more per transaction.[14]

Panel B reports the IV estimates. The first stage is strong, with a Kleibergen–Paap rk Wald F-statistic of 207.80.[15] The IV estimates point to the same conclusions, though the point estimates are larger than their matched-DID counterparts, consistent with the IV identifying effects for firms whose adoption responds to related-market diffusion rather than for the average adopter (Section 3.3). Section 4.3 shows that these results are robust to alternative estimators, samples, and outcome definitions.

[13]For outcomes transformed using $\log(x + 1)$, percentage effects are calculated as $100 \times [\exp(\hat{\beta}) - 1]$.

[14]One caveat on the price margin: because an order may contain multiple units of a product, average order price mixes prices, discounts, and quantities purchased, so its stability does not completely rule out pricing responses.

[15]The IV specifications use the unstacked firm–month panel, so their observation counts differ from those of the matched-DID specifications.

4.2 Business Capabilities

We next examine whether AI adoption strengthens firms' broader operating capabilities along three dimensions: reaching customers, adapting product offerings, and maintaining quality. While AI can directly perform certain tasks, improvements in these dimensions also depend on firms translating AI-provided information and guidance into actual operating decisions and activities.

Reaching Customers In online markets, firms must attract potential customers and convert their attention into sales. AI may help by improving how products are presented and promoted. Event studies in Appendix Figure D.1 show that adopters reach and acquire customers more effectively, with gains persisting throughout the post-adoption period. In Panel A of Table III, adopters' page views, new visitors, and new customers rise by roughly 45-50%. The new-customer conversion rate rises by 7.2% and new-customer sales by 75.6%, indicating that the additional attention translates into purchases rather than idle traffic.

One institutional detail matters for interpretation: adoption status is not an input into the platform's product-ranking or traffic-allocation algorithms, so the reach effect is not a mechanical traffic bonus for adopters. Instead, the reach effects are consistent with improvements in product presentation and operating performance that make adopters' products more attractive to consumers.

Improving Product Offerings Firm growth also depends on what firms sell. Platform AI may help on two margins: lowering the cost of introducing new products, and identifying new-product opportunities through providing information about demand, competitors, and potential products. Yet identifying opportunities does not guarantee action: firms must still decide which products to introduce and follow through with actual launches. Event studies in Appendix Figure D.2 show persistent increases in new-product activity after adoption. Panel A of Table IV reports positive effects on product introductions and category expansion. Adopters introduce 66.2% more new products and enter 62.6% more product categories, deepening and broadening their offerings. The new launches also perform better: the new-product success rate rises by 10.4%, and sales from

newly introduced products increase by 50.5%. Together, these results indicate greater product experimentation and stronger sales performance of newly introduced products.

Maintaining Quality Rapid growth can strain quality: more customers, more new products, and more orders all raise the burden on service and fulfillment. Platform AI may help firms carry this load by improving product listings and supporting customer service, fulfillment, and after-sales operations, although maintaining quality still depends on how effectively firms use these capabilities.

Event studies in Appendix Figure D.3 present the dynamic effects on quality outcomes, and Panel A of Table V reports the average effects. Product quality improves: the refund rate attributed to firms falls by 2.7%, and the negative review rate falls by 21.9%. Service quality also improves: the firm-liability dispute rate falls by 20.6%, delivery time falls by 1.5%, and the 24-hour pickup on-time rate rises by 1.9%, in line with the operational guidance, shipment monitoring, and after-sales support provided by the business-operation tool. The one outcome that moves the other way is the overall refund rate, which rises by 4.2%. The increase comes entirely from refunds attributed to consumers (for example, change of mind or ordering errors), which rise by 5.4%, consistent with broader reach and more product experimentation creating additional consumer-product mismatch.[16] In short, adopters improve product and service quality along every margin that firms control, even as their expanding reach brings more consumer-driven returns.

Panel B of Tables III–V reports the corresponding IV estimates, which are consistent with the matched-DID patterns across all three capability margins. Taken together, these findings show that AI adoption helps firms reach more customers, adapt their product offerings, and improve product and service quality as they expand. Many of these improvements depend on firms' own operating decisions and implementation, rather than AI-generated outputs alone. The evidence thus suggests that task-level AI assistance can translate into broader firm-level gains by strengthening firms' effective operating capabilities.

[16]The remaining refunds form a residual category (platform promotions, negotiated cancellations, third-party logistics issues, and administrative reasons) that bears little on product quality or firm performance, so we do not analyze them.

4.3 Robustness Analyses

We conduct a host of robustness analyses and check the baseline results along three dimensions: different estimation designs, the sample, and the outcome definition. Appendix D reports the details.

Alternative Estimators The IV estimates are robust to province-by-month and category-by-month fixed effects and to time-varying controls for baseline firm characteristics (Appendix Table C.1). For the five outcomes with many zeros that we estimate by PPML in the baseline, $\log(1+x)$ estimates are qualitatively similar: effects have the same sign and are statistically significant in both the DID and IV specifications (Appendix Table D.1).[17]

The stacked matched design avoids the weighting problems of two-way fixed-effects estimators under staggered adoption, but it relies on never-adopters as controls. We therefore re-estimate the main effects with two alternatives. The estimator of Borusyak, Jaravel, and Spiess (2024) uses the untreated observations of adopters as well as never-adopters to impute counterfactual outcomes. It implies a 40.1% increase in sales, with effects of the same sign and significance on orders, page views, new products, and quality (Appendix Table D.2). Synthetic difference-in-differences (Arkhangelsky et al., 2021) instead reweights never-adopters within cells defined by product category, size quartile, and adoption cohort and reports similar findings (Appendix Figure D.4).

Alternative Samples The baseline sample requires positive sales in every month. Relaxing this requirement adds many smaller firms with intermittent sales. In this sample, AI adoption raises sales by 75.1%, which now also reflects more months with positive sales (Appendix Table D.3).

The baseline treatment requires adopting both functionalities. Firms that adopt only one also gain, but less: sales rise by 25.2% for Content Generation alone and 22.8% for Business Operation alone, compared with 51.6% for joint adopters (Appendix Table D.4). The two functionalities also differ in the aspects of firm performance they affect. Business Operation improves 24-hour pickup

[17] Their magnitudes are not directly comparable with the PPML percentage effects, because with zeros in the outcome, $\log(1+x)$ coefficients depend on the units of measurement (Chen and Roth, 2024).

performance and reduces firm-liability disputes by more, consistent with its role in fulfillment and after-sales service, while Content Generation has no effect on pickup performance. These patterns suggest that the two functionalities are complements and reinforce each other.

Alternative Outcome Definition Finally, the sales gains are not a mechanical by-product of new-product launches. Excluding sales from products introduced in the current month, sales still rise by 43.2% (Appendix Table D.5).

4.4 Competitive Spillovers and Market-Level Effects

Competitive Spillovers Given the large gains in sales and capabilities, a natural concern is that adopters gain at the expense of non-adopters. The historical counterfactual provides the most direct test of control-group contamination. Figure III shows that estimates using contemporaneous controls are statistically indistinguishable from those using controls' adjusted 2023 outcomes, providing no evidence that competitive spillovers inflate the baseline effects.

The other two tests point in the same direction (Appendix Tables D.6 and D.7). Category-level AI adoption is not significantly associated with exit rates, either overall or among non-adopters, across alternative market definitions and fixed-effect structures, providing little evidence of displacement operating through firm exit. Nor does the treated-control gap widen with greater competitive exposure. Together, these results make it unlikely that the firm-level effects are driven by non-adopters losing sales to adopters.

Market Expansion Figure IV reports the continuous-exposure event-study estimates. For all five market-expansion outcomes, the pre-period coefficients are small and statistically indistinguishable from zero. After AI diffusion begins, high-friction categories see larger increases in total sales and incumbent sales, so entry alone does not drive the market growth. The same categories also see more firm entry, greater entrant sales, and more new-product introductions.

Taken together, these findings point to market expansion on both the intensive and extensive

margins. Incumbent sales grow faster in higher-friction categories, and so do entry, entrant sales, and new-product introductions. Because total category sales rise as well, the pattern is difficult to reconcile with existing firms redistributing a fixed pool of demand. The three spillover tests point the same way: they provide little evidence that adopters gain at the expense of non-adopters. The evidence is therefore consistent with Prediction 4. Platform AI appears to expand activity on the platform through incumbent growth, new-firm entry, and new products, rather than merely shifting sales toward adopters.

We also examine category-level concentration. The post-launch HHI coefficients are generally negative, indicating lower concentration in categories with greater baseline frictions. However, because HHI exhibits differential pre-trends, we view this evidence as suggestive.

5 Who Gains: Heterogeneity by Firm Size

The gains from platform AI may not be evenly distributed. We examine this size gradient first for sales, then for the business capabilities that support firm growth. Appendix E reports more details.

5.1 Sales Gains by Firm Size

Because the direction of the size gradient is ambiguous ex ante (Section 2.2), we estimate treatment effects by quartile of pre-launch annual sales. Figure V shows a clear gradient. Sales rise by 61.5% in the bottom quartile and by 10.3% in the top quartile, with the two middle quartiles in between and statistically similar to each other. The proportional gains from platform AI are strongly tilted toward initially smaller firms.

The sources of growth also differ by size. Larger firms grow through more orders alone: in the top quartile, orders rise by 14.3% while average order price falls by 3.9%. The smallest firms grow on both margins: in the bottom quartile, orders rise by 53.1% and average order price by 3.1%. This second margin is not what one might expect. If AI mainly brought small firms marginal, low-value buyers, average spending would fall. Instead, among newly acquired customers, spending

per customer rises by 6.0% and the average price of their orders by 3.6% in the bottom quartile, compared with an insignificant 0.6% and a decline of 2.0% in the top quartile (Appendix Figure E.1). AI thus helps the smallest firms attract not only more customers but higher-value ones.

The size gradient remains robust across the specifications discussed in Section 4.3. Sales gains remain largest for the smallest firms in IV estimates (Appendix Table E.1), under synthetic and staggered DID estimators (Appendix Figure D.4 and Appendix Table E.2), and among single-function adopters (Appendix Figure E.2) and firms in the expanded sample allowing zero-sales months (Appendix Figure E.3). Additionally, the sample that includes zero-sales months reveals an extensive-margin effect: AI adoption increases the probability of positive sales in a given month for firms in all but the largest size quartile (Appendix Figure E.3).

A back-of-the-envelope calculation gauges the aggregate magnitude. Consider a hypothetical scenario in which firms in our baseline sample (continuing physical-goods firms with annual sales of at most RMB 5 million) adopt AI. Applying the quartile-specific percentage sales gains in Figure V to these firms' sales over the study period and aggregating yields an increase equal to roughly 2.1% of total platform sales over the same period. This is economically meaningful given the platform's scale, especially since the calculation excludes larger firms above the RMB 5 million threshold that may also benefit. We interpret this as a benchmark for economic magnitude rather than a causal or welfare effect, since it does not account for equilibrium responses across firms or markets.

5.2 Capability Gains by Firm Size

The capability improvements also concentrate among initially smaller firms, though not uniformly. In customer reach, page views and new visitors rise by 55.3% and 53.3% in the bottom quartile but by 8.9% and 8.1% in the top quartile, and new-customer sales by 81.0% versus 21.7% (Appendix Figure E.4). The new-customer conversion rate, by contrast, rises by a similar 4–6% in every quartile, so smaller firms gain mainly by drawing more visitors rather than by converting them at higher rates. Product offerings show the steepest gradient (Appendix Figure E.5): bottom-quartile firms introduce 113.2% more new products and earn 191.0% more from them, compared with 27.0% and

8.6% in the top quartile. Quality improvements are likewise concentrated at the bottom (Appendix Figure E.6). Among the smallest firms, negative reviews fall by 42.5% and firm-liability disputes by 19.3%; among the largest, neither changes significantly. Firm-attributed refunds fall by similar small amounts across all quartiles. The IV estimates (Appendix Tables E.3–E.5) are consistent with these patterns.

In terms of the framework in Section 2.2, this pattern indicates that duplication dominates complementarity. The same tools fill larger gaps for firms with less precise information and more limited operating capabilities, and any advantage larger firms have in implementing AI output is not large enough to offset this. Section 6 examines the two channels behind this gradient.

6 Mechanisms

Two mechanisms, identified in Predictions 2 and 3, help explain both the average gains and their tilt toward smaller firms: platform AI lowers the cost of routine work, and it reduces information frictions. We examine each in turn.

6.1 Operating-Cost Channel

Routine business activities require labor, time, and specialized expertise. By making centrally developed capabilities available to every firm, platform AI can lower the cost of these routine tasks, allowing firms to expand these activities while enhancing labor efficiency.

Direct Evidence on Operating Activities We first examine whether AI adoption expands routine operating activities while achieving higher labor efficiency. Table VI shows both margins. Adopters produce 62.6% more images, with images per new and per existing product rising by 3.8% and 1.4%, and they update listings more often, with modifications per product up 124.8%. At the same time, human customer-service time per order falls by 5.6%.[18] The IV estimates in Panel B are

[18]Higher customer-service efficiency may allow workers to focus more on complex tasks. Interviews with the platform's business teams, conducted in March and August 2025 and October 2026, suggest that some adopting firms

qualitatively similar.

Appendix Figure E.7 shows that these changes are larger for smaller firms, which raise content production by more and achieve higher labor efficiency. This is the pattern the cost channel predicts if routine activities weigh more heavily on smaller firms. The benchmarks below quantify that burden.

Cost-Equivalent Benchmarks To gauge the economic magnitude of the activities supported by platform AI, we construct two cost-equivalent benchmarks for obtaining comparable capabilities from external providers (see Appendix F for details). The first values conventional image and promotional-video production using survey-based unit prices. The second applies a representative fixed-fee-plus-commission pricing structure for outsourced e-commerce operating support. For each benchmark, we report both a baseline calculation and a growth-adjusted calculation that incorporates estimated responses to AI adoption. These measures represent the market-price equivalent of comparable activities, not firms' observed expenditures or estimates of realized cost savings.

Table VII reports both benchmarks by size quartile for all platform firms satisfying the baseline sample restrictions, and both show a steep size gradient. The median content benchmark equals 22.7% of annual gross merchandise value (GMV) for bottom-quartile firms but 0.4% for top-quartile firms (Panel A). The median outsourced-support benchmark is roughly three times annual GMV at the bottom but 4.9% at the top (Panel B): fixed service fees put comparable support out of reach at small scale.

Heterogeneity by Baseline Operating-Cost Burden To test whether the cost channel accounts for the size gradient, we construct three firm-level burden measures from the benchmarks, each scaled by pre-launch sales: *Content Input*, the implied cost of baseline content production; *Service Input*, the implied cost of comparable outsourced support; and *Combined Input*, their sum. As an independent, region-based measure of the cost of labor-intensive tasks, we also use the 2023

shift workers from routine customer service toward more judgment-intensive tasks and new roles, such as product management, identifying customer needs, and training AI customer-service agents.

average wage of urban private wholesale and retail employees in the firm's province.

Table VIII reports the results. All four interactions are positive and statistically significant (the wage interaction at the 10% level in the DID specification): sales gains are proportionally larger for firms with heavier cost burdens and for firms in higher-wage provinces. The IV estimates show the same pattern.[19] Because these burdens fall disproportionately on smaller firms, the results tie the cost channel directly to the size gradient in Section 5. Together with the cost benchmarks in Table VII, the evidence supports the operating-cost channel: platform AI lowers the cost of routine work, and the gains are largest for the firms for which comparable capabilities would otherwise have been most costly to obtain externally.

6.2 Information-Friction Channel

Running an online business requires knowing what consumers want, which products fit the market, and how the firm's own operations are performing — knowledge that larger firms assemble with dedicated staff and smaller firms often do without. By combining each firm's data with platform-wide signals and turning both into guidance, platform AI can reduce these frictions, improving firms' decisions and the information they provide to consumers.

Direct Evidence on Information Provision Adopters' listings become substantially more informative. In Table IX, product titles lengthen by 20.2%, and attribute coverage rises by 28.4% for existing products and 32.9% for new ones. Richer titles may improve product discoverability, while more complete attributes can help consumers assess product fit (Dimoka, Hong, and Pavlou, 2012; Hirsch et al., 2025). In our baseline data, firms' average product-title length is positively correlated with firm size, firm age, clicks, and page views. The increase for existing products is especially telling: these products were already listed, so the added attributes are not a mechanical by-product of introducing new products; rather, firms appear to refine their listings as they learn

[19]Under the assumed common fixed fee and commission rate, standardized Service Input is equivalent to standardized inverse baseline GMV. The interaction with Combined Input remains positive and statistically significant after controlling for treatment-effect heterogeneity associated with inverse baseline GMV.

which characteristics and keywords matter to consumers. Moreover, unlike AI-generated titles and images, structured attributes such as product materials and specifications require firms to verify product characteristics and supply the relevant information themselves. These findings suggest that AI adoption improves information disclosure through firms' own adjustments to their listings, beyond AI-generated content alone.

Heterogeneity by Baseline Information Frictions We next test whether AI generates larger gains in markets with more severe pre-existing information frictions. We construct six predetermined measures at the Level-3 category level to capture differences across narrowly defined product markets.[20] These measures capture two dimensions of information frictions: matching quality (conversion rates and refund rates) and market instability (firm exit, sales volatility, product survival, and product turnover). Appendix Table A.3 defines each measure. Table X shows a consistent pattern across all six: firms gain more from AI in markets where matching quality is poorer and demand is harder to predict.

Appendix Figure E.8 shows that improvements in product information provision are widespread across firm sizes. Importantly, Appendix Figure E.9 shows that sales gains from AI adoption rise more steeply with baseline information frictions among smaller firms, particularly those in the bottom quartile. This pattern is consistent with platform-wide information being most valuable to firms with the least information of their own.

Together, the evidence supports the information-friction channel. Platform AI enriches the information firms provide and use, and the gains are largest where information was scarcest: in markets with poor matching and unpredictable demand, and for smaller firms.

[20]We use Level-3 categories because firms often operate across multiple Level-4 categories, making assignment to a single detailed category less representative of their primary business. Level-3 provides a broader classification of firms' main activities while retaining substantially more cross-category variation than Level-2.

7 Conclusion

Generative AI's task-level gains matter for the broader economy only if they translate into firm growth. Using complete records of AI use linked to firms' monthly operations, we show that they do. Adoption raises sales by about half, mainly through more orders rather than higher prices, and strengthens firms' business capabilities that underlie growth: reaching customers, adapting product offerings, and maintaining quality. Lower operating costs and reduced information frictions help explain these gains. At the market level, categories where AI matters more see faster growth in sales, entry, and new products, consistent with market expansion rather than pure reallocation.

The gains are largest, in proportional terms, for the smallest firms. Developing AI involves strong economies of scale, and the complementary investments needed to use it have so far favored large firms. In our setting, the platform makes these investments centrally and delivers ready-to-use applications inside the software firms already use, and the firms with the fewest resources of their own gain the most. Scale economies in producing AI need not become scale advantages in using it.

Several longer-run questions remain open. During our sample period, adopters reduce customer-service staff time per order while simultaneously increasing sales by about half. Over a longer horizon, the efficiency effect of AI may allow firms to grow without proportionate increases in hiring, or the growth it generates may lead firms to hire more workers and shift them from routine tasks toward judgment-intensive work. Another question is whether the operating gains we document build lasting human and organizational capital as workers learn from AI outputs and recommendations. The answers will show whether these gains persist and how they are shared between firms and workers.

Funding

Lu Fang acknowledges financial support from the National Natural Science Foundation of China [Grants 72501258 and 72192803], the National Social Science Fund of China [Grant 22&ZD081]. Zhe Yuan acknowledges financial support from the National Key Research and Development Project

[Grant 2024YFB3312900], the National Natural Science Foundation of China [Grants 72673168 and 72141305], Major Project of the National Social Science Fund of China [Grant 25&ZD149 and 26ZDA030].

References

Acemoglu, Daron, Gary W. Anderson, David N. Beede, Catherine Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas, "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey," in Basu, Susanto, Lucy Eldridge, John Haltiwanger, and Erich Strassner eds. *Technology, Productivity, and Economic Growth*, 83 of Studies in Income and Wealth, 13–56, Chicago: University of Chicago Press (2025).

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

——— "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33 (2019), 3–30.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska, "The Demand for AI Skills in the Labor Market," *Labour Economics*, 71 (2021), 102002.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic Difference-in-Differences," *American Economic Review*, 111 (2021), 4088–4118.

Athey, Susan, Dean Eckles, and Guido W. Imbens, "Exact P-Values for Network Interference," *Journal of the American Statistical Association*, 113 (2018), 230–240.

Athey, Susan and Fiona M. Scott Morton, "Artificial Intelligence, Competition, and Welfare," in Agrawal, Ajay K., Erik Brynjolfsson, and Anton Korinek eds. *The Economics of Transformative AI*, Chap. 4, 95–114, Chicago: University of Chicago Press (2026).

Babina, Tania, "Understanding Firms' AI Efforts and Their Economic Impact," *The Review of Corporate Finance Studies*, 15 (2026), 967–1013.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 151 (2024), 103745.

Bäck, Asta, Arash Hajikhani, Angela Jäger, Torben Schubert, and Arho Suominen, "Return of the Solow-paradox in AI? AI-adoption and firm productivity," Papers in Innovation Studies 2022/1, Lund University, CIRCLE–Centre for Innovation Research, 2022.

Barwick, Panle Jia, Siyu Chen, Chao Fu, and Teng Li, "Digital Distractions with Peer Influence: The Impact of Mobile App Usage on Academic and Labor Market Outcomes," *The Quarterly Journal of Economics*, 141 (2026), 1–49.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81 (2013), 1347–1393.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89 (2022), 181–213.

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

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, 117 (2002), 339–376.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," *The Quarterly Journal of Economics*, 140 (2025), 889–942.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 13 (2021), 333–372.

Cai, Junhui Jeff, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao, and Wu Zhu, "AI as "Co-Founder": GenAI for Entrepreneurship," arXiv preprint arXiv:2512.06506, 2025.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," *American Economic Review* (2026), Forthcoming.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 134 (2019), 1405–1454.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros? Some Problems and Solutions," *The Quarterly Journal of Economics*, 139 (2024), 891–936.

Cohen, Wesley M. and Daniel A. Levinthal, "Absorptive Capacity: A New Perspective on Learning and

Innovation," *Administrative Science Quarterly*, 35 (1990), 128–152.

Comin, Diego and Thomas Philippon, "The Rise in Firm-Level Volatility: Causes and Consequences," *NBER Macroeconomics Annual*, 20 (2005), 167–201.

Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz, "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers," *Management Science* (2026), Articles in Advance.

Czarnitzki, Dirk, Gastón P Fernández, and Christian Rammer, "Artificial Intelligence and Firm-Level Productivity," *Journal of Economic Behavior & Organization*, 211 (2023), 188–205.

Damioli, Giacomo, Vincent Van Roy, and Daniel Vertesy, "The Impact of Artificial Intelligence on Labor Productivity," *Eurasian Business Review*, 11 (2021), 1–25.

Demirer, Mert, Andrey Fradkin, and Nadav Tadelis, "The Emerging Market for Intelligence: How Firms Buy and Sell AI," *Journal of Economic Perspectives*, 40 (2026), 23–46.

DeStefano, Timothy, Richard Kneller, and Jonathan Timmis, "Cloud Computing and Firm Growth," *The Review of Economics and Statistics*, 107 (2025), 1638–1651.

Dimoka, Angelika, Yili Hong, and Paul A. Pavlou, "On Product Uncertainty in Online Markets: Theory and Evidence," *MIS Quarterly*, 36 (2012), 395–426.

Duflo, Esther, "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment," *American Economic Review*, 91 (2001), 795–813.

Eisfeldt, Andrea L., Gregor Schubert, Miao Ben Zhang, and Bledi Taska, "Generative AI and Firm Values," *The Journal of Finance* (2026), Early View.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108 (2000), 874–904.

Goldberg, Samuel and H Tai Lam, "Generative AI in Equilibrium: Evidence from a Creative Goods Marketplace," (2025), Stanford Graduate School of Business Working Paper No. 4265.

Goldfarb, Avi and Daniel Trefler, "Artificial Intelligence and International Trade," in Agrawal, Ajay, Joshua Gans, and Avi Goldfarb eds. *The Economics of Artificial Intelligence: An Agenda*, Chap. 19, 463–492, Chicago: University of Chicago Press (2019).

Hagiu, Andrei and Julian Wright, "Artificial Intelligence and Competition Policy," *International Journal of Industrial Organization*, 103 (2025), 103134.

Helm, Ines, "National Industry Trade Shocks, Local Labour Markets, and Agglomeration Spillovers," *The Review of Economic Studies*, 87 (2020), 1399–1431.

Hirsch, Sharon, Ido Guy, Slava Novgorodov, Gal Lavee, and Bracha Shapira, "What Is in a Title? Characterizing Product Titles in E-Commerce," *Expert Systems with Applications*, 287 (2025), 127702.

Horn, Danea, Adam Sacarny, and R. Annetta Zhou, "Technology Adoption and Market Allocation: The Case of Robotic Surgery," *Journal of Health Economics*, 86 (2022), 102672.

Jovanovic, Boyan and Peter L. Rousseau, "AI and Task Efficiency," *Journal of Monetary Economics*, 157 (2026), 103877.

Korinek, Anton and Jai Vipra, "Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence," *Economic Policy*, 40 (2025), 225–256.

Krakowski, Sebastian, Darek Haftor, Johannes Luger, Natallia Pashkevich, and Sebastian Raisch, "Human-Centered Artificial Intelligence: A Field Experiment," *Management Science*, 72 (2026), 57–72.

Lu, Yao, Gordon M. Phillips, and Jia Yang, "The Impact of Cloud Computing and AI on Industry Dynamics and Concentration," SSRN Working Paper, revised December 24, 2025, 2024.

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia S. Foster, and Nikolas Zolas, "AI Adoption in America: Who, What, and Where," *Journal of Economics & Management Strategy*, 33 (2024), 375–415.

McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson, "The Rise of Industrial AI in America: Microfoundations of the Productivity J-Curve(s)," CES Working Paper 25-27, U.S. Census Bureau, Center for Economic Studies, 2025.

Milgrom, Paul R and Steven Tadelis, "How Artificial Intelligence and Machine Learning Can Impact Market Design," in Agrawal, Ajay, Joshua Gans, and Avi Goldfarb eds. *The Economics of Artificial Intelligence: An Agenda*, 567–586, Chicago: University of Chicago Press (2019).

Noy, Shakked and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381 (2023), 187–192.

Nunn, Nathan and Nancy Qian, "US Food Aid and Civil Conflict," *American Economic Review*, 104 (2014), 1630–1666.

Otis, Nicholas G., Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand M. Koning, "The Uneven Impact of Generative Artificial Intelligence on Entrepreneurial Performance: Evidence from a Field

Experiment in Kenya," *Management Science* (2026), Articles in Advance.

da Silva Marioni, Larissa, Ana Rincon-Aznar, and Francesco Venturini, "Productivity Performance, Distance to Frontier and AI Innovation: Firm-Level Evidence from Europe," *Journal of Economic Behavior & Organization*, 228 (2024), 106762.

Varian, Hal R., "Artificial Intelligence, Economics, and Industrial Organization," in Agrawal, Ajay, Joshua Gans, and Avi Goldfarb eds. *The Economics of Artificial Intelligence: An Agenda*, Chap. 16, 399–419, Chicago: University of Chicago Press (2019).

Wing, Coady, Seth M. Freedman, and Alex Hollingsworth, "Stacked Difference-in-Differences," NBER Working Paper 32054, National Bureau of Economic Research, Revised August 2026, 2024.

Table I
Summary Statistics of Main Variables

Variable	Mean	P5	P50	P95
Panel A. $\log(1+x)$*-transformed monthly firm outcomes*				
Sales	8.08	5.38	8.18	10.76
Orders	3.69	1.39	3.56	6.31
Average Order Price	4.46	2.61	4.40	6.85
Firm Page Views	7.59	4.84	7.57	10.41
New Visitors	6.61	3.85	6.58	9.39
New Customers	3.14	0.69	3.00	5.88
New-Customer Sales	7.25	3.33	7.50	10.36
Number of New Categories	0.10	0.00	0.00	0.69
Number of New Products	0.58	0.00	0.00	3.14
New-Product Sales	1.18	0.00	0.00	7.76
Number of Images	1.21	0.00	0.00	5.04
Modifications per Product	0.24	0.00	0.00	1.16
Human Customer-Service Time per Order	4.72	0.00	5.02	7.21
Product-Title Length	2.94	2.29	2.86	3.50
Delivery Time	4.15	3.72	4.11	4.66
Panel B. Monthly firm rates and shares relative to control-group statistics				
Refund Rate	1.210	–	1.938	1.200
Product Negative Review Rate	0.855	–	–	0.829
Firm-Liability Dispute Rate	0.787	–	–	0.612
New-Product Success Rate	1.007	–	0.831	1.000
New-Customer Conversion Rate	0.985	1.594	1.027	0.979
24-Hour Pickup On-Time Rate	1.037	1.364	1.054	1.000

Notes: This table reports summary statistics for the baseline matched sample used in the firm-level analysis. Panel A reports monthly firm outcomes transformed using $\log(1+x)$. To comply with the platform's disclosure requirements, Panel B reports rate and share outcomes relative to the corresponding statistics among matched control firms: each reported mean or percentile is divided by the corresponding control-group mean or percentile. A value of one therefore denotes equality with the control-group benchmark; a dash indicates that the corresponding control-group statistic is zero and the ratio is undefined. Table A.1 provides variable definitions.

Table II
Average Effects of AI Adoption on Firm Sales Performance

Variables	(1) Sales	(2) Orders	(3) Average Order Price
Panel A. Difference-in-differences estimates			
Treat × *Post*	0.416***	0.385***	0.009
	(0.0127)	(0.0134)	(0.00656)
Firm-Cohort FE	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓
% Change	51.6%	47.0%	–
Observations	6,020,256	6,020,256	6,020,256
R-squared	0.664	0.711	0.817
Panel B. Instrumental-variable estimates			
AI Adoption	0.733***	0.687***	0.00551
	(0.0456)	(0.0425)	(0.0133)
Firm FE	✓	✓	✓
Month FE	✓	✓	✓
First-stage	2.255***	2.255***	2.255***
KP F-statistic	207.80	207.80	207.80
Observations	4, 974, 574	4, 974, 574	4, 974, 574

Notes: This table reports the average effects of platform AI adoption on monthly sales, orders, and average order price. Panel A reports the baseline matched-DID estimates, comparing firms that jointly adopt Content Generation and Business Operation with cohort-matched firms that never adopt either functionality. Panel B reports IV estimates on the unstacked firm-month panel, instrumenting AI Adoption with the instrument in equation (3). Panel A includes firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B includes firm and calendar-month fixed effects. All outcomes are transformed using $\log(1+x)$. The percentage effects in Panel A are calculated as $100[\exp(\hat{\beta}) - 1]$, with a dash denoting a statistically insignificant estimate. Standard errors are clustered at the Level-2 category level. Panel B also reports the first-stage coefficient and Kleibergen-Paap rk Wald F-statistic. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table III
Effects of AI Adoption on Customer Reach

Variables	(1) Firm Page Views	(2) New Visitors	(3) New Customers	(4) New-Customer Conversion Rate	(5) New-Customer Sales
Panel A. Difference-in-differences estimates					
Treat × *Post*	0.378***	0.369***	0.403***	0.0722***	0.563***
	(0.0122)	(0.0117)	(0.0133)	(0.00635)	(0.0138)
Firm-Cohort FE	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓
% Change	45.9%	44.6%	49.6%	7.2%	75.6%
Observations	6,020,256	6,020,256	6,020,256	6,019,721	6,020,256
R-squared	0.753	0.747	0.707	0.527	0.582
Panel B. Instrumental-variable estimates					
AI Adoption	0.756***	0.779***	0.730***	0.0456***	0.987***
	(0.0613)	(0.0633)	(0.0463)	(0.0159)	(0.0675)
Firm FE	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
First-stage	2.255***	2.255***	2.255***	2.254***	2.255***
KP F-statistic	207.80	207.80	207.80	207.80	207.80
Observations	4,974,574	4,974,574	4,974,574	4,974,242	4,974,574

Notes: This table reports the effects of platform AI adoption on customer reach and acquisition. Panel A reports the baseline matched-DID estimates with firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B reports IV estimates with firm and calendar-month fixed effects, instrumenting AI Adoption with the instrument in equation (3). Firm Page Views, New Visitors, New Customers, and New-Customer Sales are transformed using $\log(1+x)$, with percentage effects calculated as $100[\exp(\hat{\beta})-1]$. New-Customer Conversion Rate in column (4) is the share of first-time visitors who make their first purchase from the firm within the same month. It is observed only in firm-months with new visitors and is normalized by its matched-control pre-adoption mean, so $100\hat{\beta}$ gives the percentage change relative to that benchmark. Standard errors are clustered at the Level-2 category level. Panel B also reports the first-stage coefficient and Kleibergen-Paap rk Wald F-statistic. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table IV
Effects of AI Adoption on Product Offerings

Variables	(1) Number of New Categories	(2) Number of New Products	(3) New-Product Success Rate	(4) New-Product Sales
Panel A. Difference-in-differences estimates				
Treat × *Post*	0.486***	0.508***	0.104***	0.409***
	(0.0208)	(0.0175)	(0.00680)	(0.0224)
Firm-Cohort FE	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓
% Change	62.6%	66.2%	10.4%	50.5%
Observations	4,224,902	5,208,610	1,868,231	5,080,504
R-squared	0.210	0.583	0.425	0.682
Panel B. Instrumental-variable estimates				
AI Adoption	0.943***	0.437***	0.0455*	0.972***
	(0.0560)	(0.0728)	(0.0262)	(0.0687)
Firm FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
First-stage	–	–	1.690***	–
KP F-statistic	–	–	167.16	–
Observations	3,531,814	4,350,324	1,509,819	4,251,456

Notes: This table reports the effects of platform AI adoption on product offerings. Panel A reports matched-DID estimates with firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B reports IV estimates with firm and calendar-month fixed effects, using the instrument in equation (3). To accommodate outcomes with a substantial mass at zero and avoid the unit sensitivity of $\log(1+x)$ transformations, columns (1), (2), and (4) use PPML in Panel A and IV-PPML in Panel B, with percentage effects calculated as $100[\exp(\hat{\beta})-1]$. New-Product Success Rate in column (3) is the share of newly introduced products generating at least one sale, observed only in firm-months with new introductions and normalized by the matched-control pre-adoption mean, with percentage effects calculated as $100\hat{\beta}$. PPML excludes separated observations and singleton fixed-effect groups where applicable. First-stage statistics are reported only for the linear IV specification, as IV-PPML does not yield directly comparable conventional first-stage statistics. Corresponding $\log(1+x)$ IV-2SLS estimates are reported in Appendix Table D.1. Standard errors are clustered at the Level-2 category level. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

Table V
Effects of AI Adoption on Quality

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Product Quality				Service Quality		
	Refund Rate	Consumer-Reason Refund Rate	Firm-Reason Refund Rate	Product Negative Review Rate	Firm-Liability Dispute Rate	24-Hour Pickup On-Time Rate	Delivery Time
Panel A. Difference-in-differences estimates							
Treat × *Post*	0.0420***	0.0539***	−0.0265***	−0.219***	−0.206***	0.0185***	−0.0150***
	(0.0142)	(0.00629)	(0.00850)	(0.0291)	(0.0308)	(0.00142)	(0.00169)
Firm-Cohort FE	✓	✓	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓	✓	✓
% Change	4.2%	5.4%	−2.7%	−21.9%	−20.6%	1.9%	−1.5%
Observations	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256
R-squared	0.381	0.385	0.246	0.506	0.187	0.475	0.586
Panel B. Instrumental-variable estimates							
AI Adoption	0.219**	0.327**	−0.213***	−0.725*	−0.376***	0.0375***	−0.0251**
	(0.111)	(0.128)	(0.0552)	(0.391)	(0.105)	(0.00151)	(0.0100)
Firm FE	✓	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓	✓
First-stage	2.255***	2.255***	2.255***	2.255***	2.255***	2.255***	2.255***
KP F-statistic	207.80	207.80	207.80	207.80	207.80	207.80	207.80
Observations	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574

Notes: This table reports the effects of platform AI adoption on product and service quality. Panel A reports the baseline matched-DID estimates with firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B reports IV estimates with firm and calendar-month fixed effects, instrumenting AI Adoption with the instrument in equation (3). The rate outcomes in columns (1)-(6) are normalized by their matched-control pre-adoption means, so $100\hat{\beta}$ gives the percentage change relative to the corresponding benchmark. Delivery Time in column (7) is transformed using $\log(1+x)$, with percentage effects calculated as $100[\exp(\hat{\beta})-1]$. Standard errors are clustered at the Level-2 category level. Panel B also reports the first-stage coefficient and Kleibergen-Paap rk Wald F-statistic. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

Table VI
Effects of AI Adoption on Operating Activities

Variables	(1) Number of Images	(2) Images per New Product	(3) Images per Existing Product	(4) Modifications per Product	(5) Human Customer-Service Time per Order
Panel A. Difference-in-differences estimates					
Treat × *Post*	0.486***	0.0373***	0.0144***	0.810***	−0.0574***
	(0.0221)	(0.00444)	(0.00443)	(0.0359)	(0.00620)
Firm-Cohort FE	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓
% Change	62.6%	3.8%	1.4%	124.8%	−5.6%
Observations	5,602,322	1,868,231	6,020,256	5,440,314	6,020,256
R-squared	0.583	0.332	0.140	0.327	0.753
Panel B. Instrumental-variable estimates					
AI Adoption	0.585***	0.173***	0.0661***	0.350***	−0.172***
	(0.0569)	(0.0220)	(0.0109)	(0.0618)	(0.0251)
Firm FE	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
First-stage	–	1.690***	2.255***	–	2.255***
KP F-statistic	–	167.16	207.80	–	207.80
Observations	4,660,480	1,509,819	4,974,574	4,583,744	4,974,574

Notes: This table reports the effects of platform AI adoption on operating activities. Panel A reports matched-DID estimates with firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B reports IV estimates with firm and calendar-month fixed effects, using the instrument in equation (3). To accommodate outcomes with a substantial mass at zero and avoid the unit sensitivity of $\log(1+x)$ transformations, columns (1) and (4) use PPML in Panel A and IV-PPML in Panel B; the remaining columns use linear models of $\log(1+x)$. The % Change row reports $100[\exp(\hat{\beta})-1]$. Images per New Product is observed only in firm-months with new introductions. PPML excludes separated observations and singleton fixed-effect groups where applicable. Because IV-PPML has no conventional linear first stage, first-stage statistics are reported only for linear IV columns; corresponding $\log(1+x)$ IV-2SLS estimates appear in Appendix Table D.1. Standard errors are clustered at the Level-2 category level. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

Table VII
Cost-Equivalent Benchmarks for AI-Assisted Business Activities

	Baseline Calculation			Growth-Adjusted Calculation		
Firm Size Group	P25	Median	P75	P25	Median	P75
Panel A. Content-Production Benchmark						
Bottom	4.70%	22.74%	122.93%	10.02%	48.44%	261.86%
Lower-Middle	1.07%	3.61%	13.99%	1.78%	5.99%	23.23%
Upper-Middle	0.54%	1.84%	7.28%	0.70%	2.39%	9.47%
Top	0.09%	0.38%	1.63%	0.12%	0.48%	2.07%
Panel B. Outsourced-Support Benchmark						
Bottom	127.1%	299.9%	970.7%	128.9%	301.8%	972.6%
Lower-Middle	32.7%	40.5%	50.9%	33.7%	41.5%	51.9%
Upper-Middle	15.4%	18.4%	22.1%	16.5%	19.5%	23.2%
Top	3.8%	4.9%	7.4%	4.1%	5.2%	7.7%

Notes: This table reports cost-equivalent benchmarks for conventional content production (Panel A) and outsourced operating support (Panel B), expressed as shares of baseline annual GMV. The cost benchmarks are calculated using all platform firms that satisfy the baseline sample restrictions, rather than the matched estimation sample. Within quartiles of pre-launch annual sales (cutoffs: RMB 24,460, 63,088, and 147,716), we report the 25th, 50th, and 75th percentiles. The growth-adjusted calculations incorporate estimated post-adoption changes in content-production activity or sales, using PPML estimates for new-product growth in Panel A. Appendix F provides the underlying assumptions, formulas, and interpretation. The benchmarks do not measure observed expenditures or realized cost savings.

Table VIII
Heterogeneity by Baseline Operating-Cost Burden

	Difference-in-Differences				*Instrumental Variables*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	Sales							
Treatment	0.415***	0.330***	0.361***	0.417***	0.719***	0.733***	0.727***	0.706***
	(0.0126)	(0.0125)	(0.0138)	(0.0128)	(0.0246)	(0.0241)	(0.0257)	(0.0423)
Treatment × *Content Input*	0.0640**				0.158***			
	(0.0271)				(0.0334)			
Treatment × *Service Input*		0.0978***				0.0927**		
		(0.0127)				(0.0461)		
Treatment × *Combined Input*			0.124***				0.330***	
			(0.0213)				(0.0350)	
Treatment × *Local Wage Level*				0.0200*				0.0122**
				(0.0107)				(0.00549)
Observations	6,020,256	6,020,256	6,020,256	6,020,256	4,974,574	4,974,574	4,974,574	4,974,574
R-squared	0.664	0.665	0.668	0.664	–	–	–	–
KP F-statistic					459.35	447.93	469.53	506.71
Min. SW F-statistic					393.66	697.58	551.95	1,283.01

Notes: This table examines whether the sales effects of AI adoption vary with pre-adoption operating-cost burdens. The dependent variable is monthly sales, transformed using $\log(1 + x)$. Each column interacts treatment with a standardized, predetermined cost-related measure. Content Input and Service Input measure the implied costs of baseline content production and outsourced operating service, respectively, scaled by pre-launch sales; Combined Input is their sum. Local Wage Level is the 2023 provincial average wage of private-sector wholesale and retail workers. Columns (1)–(4) report matched-DID estimates (Treatment=Treat × Post), with firm-by-cohort and cohort-by-month fixed effects; columns (5)–(8) report IV estimates (Treatment=AI Adoption), with firm and calendar-month fixed effects. Treatment coefficient gives the effect at the measure's mean, while interaction coefficient captures how the effect changes with a one-standard-deviation increase. The IV specifications instrument Treatment and its interaction using the adoption instrument and its interaction with the same baseline measure. The corresponding first-stage estimates are reported in Appendix Table C.2. Relevant lower-order terms are included but not reported. Standard errors are clustered at the Level-2 category level. The bottom rows report Kleibergen-Paap rk Wald and minimum Sanderson-Windmeijer F-statistics. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table IX
Effects of AI Adoption on Product Information Provision

Variables	(1) Product-Title Length	(2) Existing-Product Attribute Coverage	(3) New-Product Attribute Coverage
Panel A. Difference-in-differences estimates			
Treat × *Post*	0.184***	0.00563***	0.00921***
	(0.00521)	(0.000631)	(0.000896)
Firm-Cohort FE	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓
% Change	20.2%	28.4%	32.9%
Observations	2,617,724	6,020,256	1,868,231
R-squared	0.398	0.302	0.417
Panel B. Instrumental-variable estimates			
AI Adoption	0.126***	0.0116***	0.000512
	(0.0172)	(0.00322)	(0.00206)
Firm FE	✓	✓	✓
Month FE	✓	✓	✓
First-stage	1.863***	2.255***	1.690***
KP F-statistic	177.28	207.80	167.16
Observations	2,272,339	4,974,574	1,509,819

Notes: This table reports the effects of platform AI adoption on information provided in product listings. Panel A reports the baseline matched-DID estimates with firm-by-cohort and cohort-by-calendar-month fixed effects; Panel B reports IV estimates with firm and calendar-month fixed effects, instrumenting AI Adoption with the instrument in equation (3). Product-Title Length in column (1) is the average number of title characters for newly introduced products or products with title modifications during the month; it is missing for firm-months with no such products. It is transformed using $\log(1+x)$, with percentage effects calculated as $100[\exp(\hat{\beta})-1]$. Existing- and New-Product Attribute Coverage in columns (2) and (3) measure the number of filled product attributes relative to the maximum observed in the same Level-2 category; column (3) is defined only for firm-months with newly introduced products. Both coverage measures are estimated in levels, with percentage effects calculated relative to their respective pre-adoption control means. Standard errors are clustered at the Level-2 category level. Panel B also reports the Kleibergen-Paap rk Wald F-statistic. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table X
Heterogeneity by Baseline Information Frictions

	Difference-in-Differences						*Instrumental Variables*					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Dependent Variable	Sales											
Treatment	0.417***	0.417***	0.415***	0.415***	0.414***	0.415***	0.545***	0.668***	0.722***	0.730***	0.685***	0.717***
	(0.0129)	(0.0104)	(0.00749)	(0.0129)	(0.00949)	(0.0137)	(0.0919)	(0.0234)	(0.0235)	(0.0460)	(0.0155)	(0.0471)
Treatment × *Conversion Rate*	−0.0224***						−0.493***					
	(0.00106)						(0.187)					
Treatment × *Refund Rate*		0.0268***						0.0982***				
		(0.00761)						(0.0183)				
Treatment × *Firm Exit Rate*			0.0225*						0.0813***			
			(0.0119)						(0.0172)			
Treatment × *Sales Volatility*				0.0288***						0.0123*		
				(0.00762)						(0.00621)		
Treatment × *Product Survival Duration*					−0.0680***						−0.138***	
					(0.0171)						(0.0133)	
Treatment × *Product Turnover Rate*						0.0275***						0.0544***
						(0.00681)						(0.00699)
Observations	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574
R-squared	0.664	0.664	0.664	0.664	0.664	0.664	–	–	–	–	–	–
KP F-statistic							126.84	403.13	508.05	104.78	684.26	105.96
Min SW F-statistic							196.74	864.87	530.15	211.10	1,285.33	231.13

Notes: This table examines whether the sales effect of AI adoption varies with pre-AI information frictions in the firm's Level-3 category. The dependent variable is monthly sales, transformed using $\log(1 + x)$. Each column interacts treatment with a standardized, predetermined category-level measure of information frictions. Lower Conversion Rate and shorter Product Survival Duration, and higher Refund Rate, Firm Exit Rate, Sales Volatility, and Product Turnover Rate indicate greater frictions. Columns (1)-(6) report matched-DID estimates (Treatment=Treat × Post), with firm-by-cohort and cohort-by-month fixed effects; columns (7)-(12) report IV estimates (Treatment=AI Adoption), with firm and calendar-month fixed effects. Treatment coefficient gives the effect at the measure's mean, while interaction coefficient captures how the effect changes with a one-standard-deviation increase. The IV specifications instrument Treatment and its interaction by the adoption instrument and its interaction with the same baseline measure. The corresponding first-stage estimates are reported in Appendix Table C.3. Relevant lower-order terms are included but not reported. Standard errors are clustered at the Level-2 category level. The bottom rows report the Kleibergen-Paap and minimum Sanderson-Windmeijer F-statistics. Negative interactions in columns (1), (5), (7), and (11) indicate larger gains where frictions are greater. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

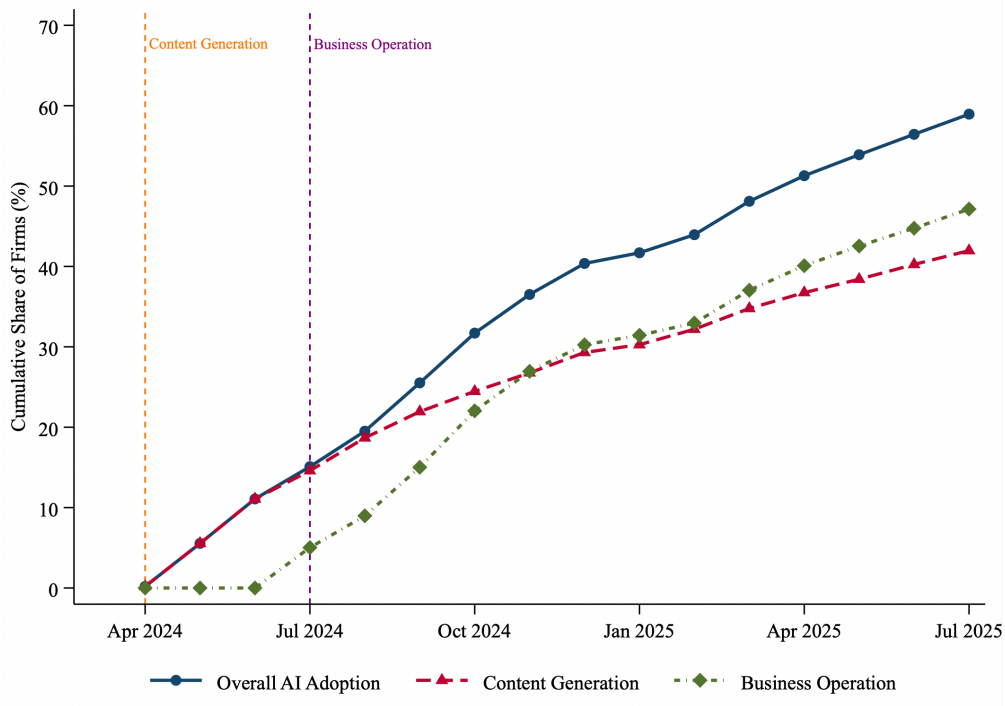
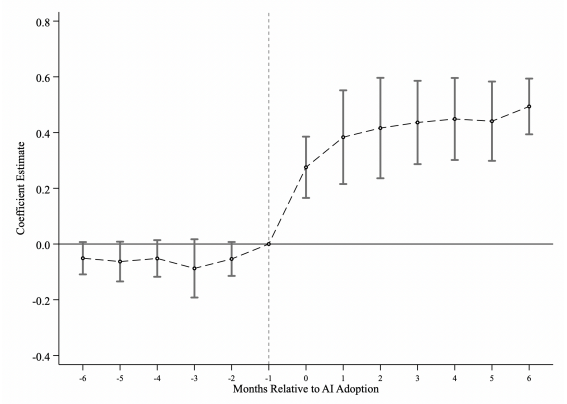
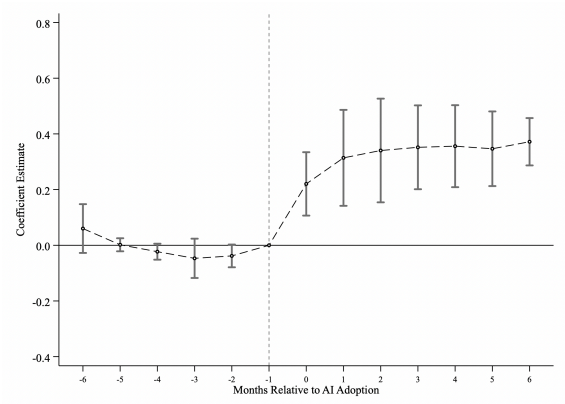


Figure I
Cumulative AI Adoption

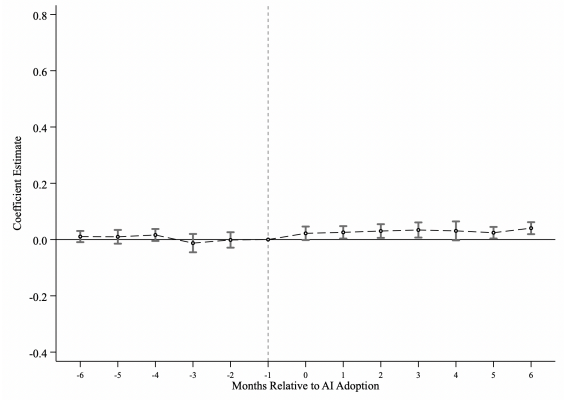
Notes: This figure plots cumulative adoption of the platform AI system among platform firms observed during the study period. A firm is counted as adopting a functionality from the first month it initiates a recorded request. The cumulative adoption rate is calculated as the number of firms that have adopted AI by each month divided by the fixed number of unique platform firms observed during the sample period. The solid line reports the share of firms adopting at least one of the two functionalities; the dashed and dash-dotted lines report the corresponding shares for Content Generation and Business Operation separately. Content Generation became available in April 2024 and Business Operation in July 2024.



(a) Sales



(b) Orders



(c) Average Order Price

Figure II
Dynamic Effects of AI Adoption on Firm Performance

Notes: This figure reports balanced event-study estimates for monthly sales, orders, and average order price, each transformed using $\log(1 + x)$. The sample contains seven adoption cohorts from July 2024 through January 2025 and uses a common six-month window before and after adoption, ensuring that all event-time coefficients are estimated from the same set of cohorts. Event time is measured relative to each treated cohort's adoption month for both treated firms and their matched controls. Month -1 is omitted, so coefficients measure changes in the treated-control gap relative to the month before adoption. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical bars denote 95% confidence intervals, with standard errors clustered at the Level-2 category level.

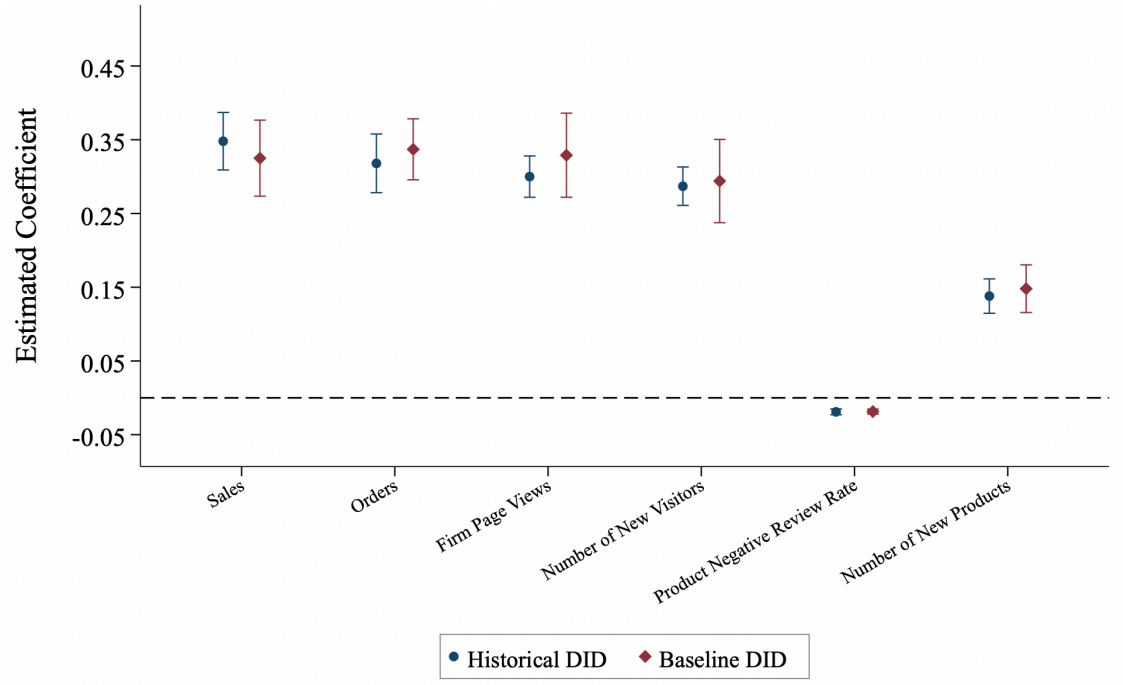
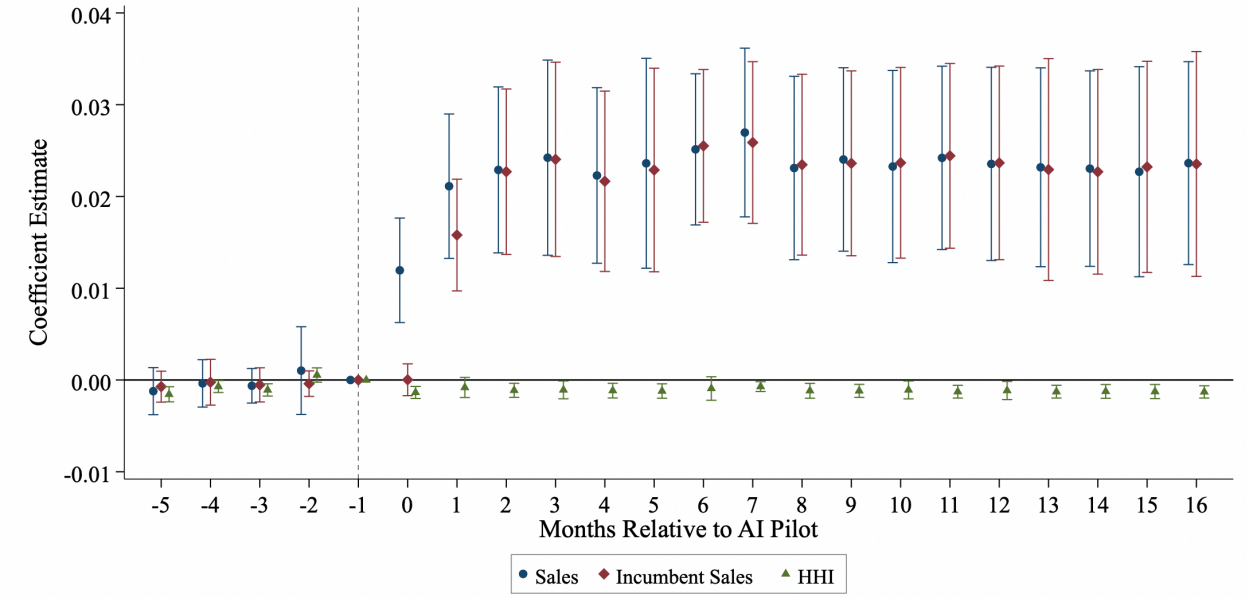


Figure III
Spillovers: Historical Counterfactual

Notes: This figure compares the baseline matched-DID estimates with a historical-counterfactual specification designed to assess potential AI-driven spillovers onto control firms. For each control firm, post-period outcomes in July–December 2024 are replaced by the firm's outcomes in the same months of 2023, adjusted using outcome-specific platform-wide growth from January–March 2023 to January–March 2024. The Product Negative Review Rate is left unadjusted because it is a rate with limited year-over-year variation. Sales, Orders, Firm Page Views, New Visitors, and Number of New Products are transformed using log(1+x), while Product Negative Review Rate is estimated in levels. Treated firms retain their observed 2024 outcomes, and we re-estimate the baseline design on the same sample. Dots report the Treat × Post coefficients, with vertical bars indicating 95% confidence intervals. All regressions include firm-by-cohort and cohort-by-calendar-month fixed effects, with standard errors clustered at the Level-2 category level.

Panel A. Aggregate and Incumbent Sales



Panel B. Entry and New Market Activity

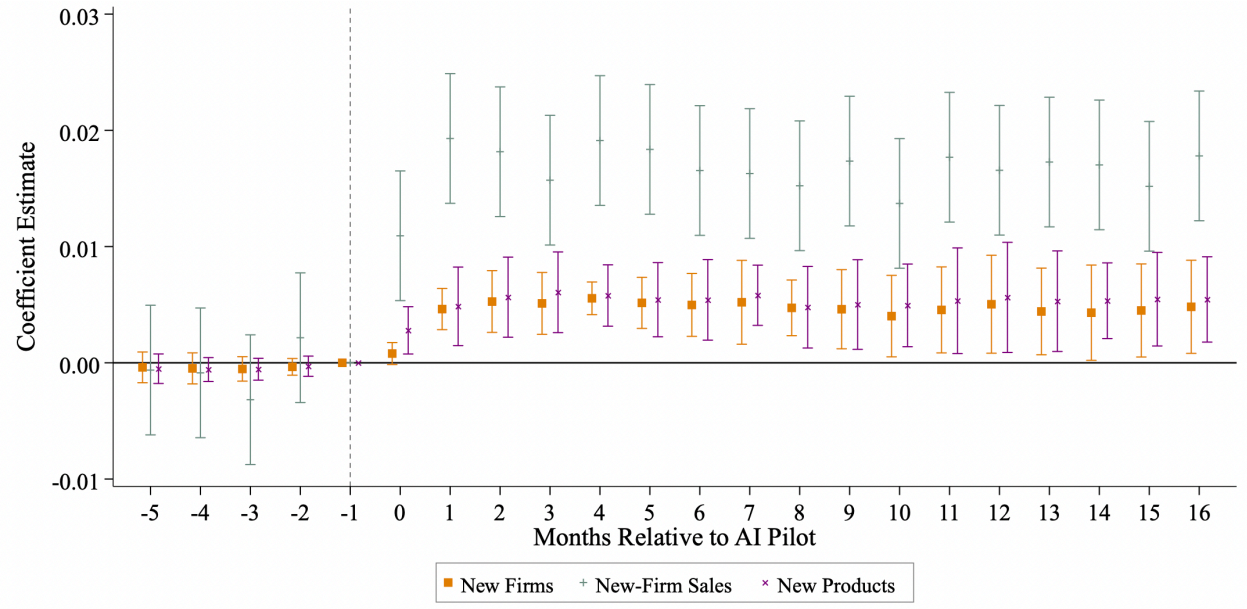


Figure IV
Market Expansion

Notes: This figure reports a category-level continuous-exposure event study of market activity around the March 2024 pilot of platform AI, using each Level-4 category's 2023 refund rate to proxy for baseline information frictions. Each point is the coefficient on the interaction between the baseline refund rate and event month indicator. Panel A reports total sales, incumbent sales, and the Herfindahl-Hirschman Index (HHI); Panel B reports firm entry, entrant sales, and new-product introductions. Firm Entry counts first-time platform entrants selling in the corresponding Level-4 category. New-Firm Sales are sales generated by these firms, while Incumbent Sales are sales generated by firms already operating on the platform. All outcomes except HHI are transformed using $\log(1+x)$; HHI is measured in levels. Event time is relative to March 2024, with February 2024 omitted. Regressions include category and calendar-month fixed effects, and vertical bars denote 95% confidence intervals based on category-clustered standard errors.

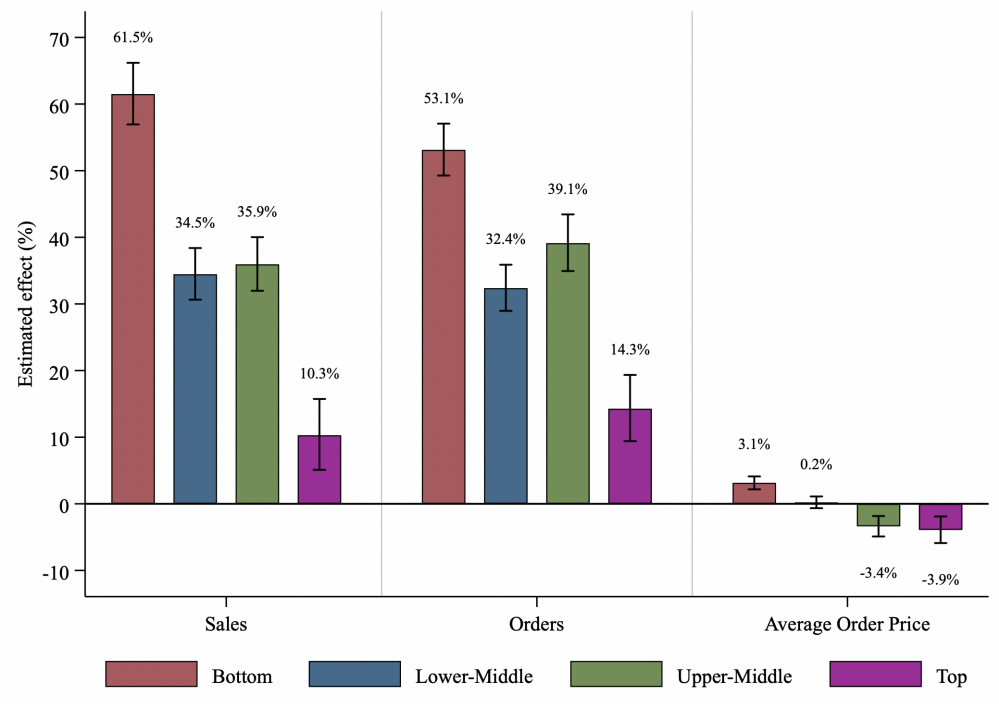


Figure V
Heterogeneous Effects of AI Adoption on Sales Performance

Notes: This figure reports matched-DID estimates of AI adoption effects on sales performance by quartile of pre-launch annual sales. The outcomes are monthly sales, orders, and average order price, each transformed using $\log(1+x)$. Bars report percentage effects calculated as $100[\exp(\hat{\beta})-1]$, and vertical lines denote 95% confidence intervals. The quartile cutoffs are RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects, with standard errors clustered at the Level-2 category level.

Appendix A Institutional Background and Data

A.1 Details of Platform AI Tool

This appendix provides illustrations of the platform AI system described in Section 2.1, covering its applications in content production and business operations, as well as how firms use its different functions.

Figure A.1 illustrates how the Content-Generation functionality supports product-content production. Panel A shows how firms can transform existing product images into listing-ready visuals by changing backgrounds, adding contextual scenes, and incorporating virtual models. Panel B shows how product images can be used to generate promotional videos featuring product demonstrations and usage scenarios. Both applications retain the focal product while reducing the need for conventional photography, video production, and editing.

Figure A.2 illustrates nine applications of Business Operation. *Operational Guidance* (Panels B–C) provides contextualized answers and actionable recommendations on platform rules and operating procedures. *Data Analytics* (Panels D–F) converts operating and market information into daily performance reports, market-trend insights, and competitor comparisons. *Market & Product Management* (Panels A, G, and I) helps firms optimize product listings, reach potential customers, identify promising new products to introduce, and find potential suppliers for sourcing them. *Service & Fulfillment* (Panel H) monitors logistics and after-sales exceptions that require firms' attention. Together, these applications illustrate how the system combines firm-specific data, platform-wide information, and operational guidance to support business decisions and routine activities. All examples are anonymized.

Figure A.3 reports the distribution of firm-initiated AI requests across the six business functions by initial firm size. *Operational Guidance* and *Content Editing* account for larger shares of requests among smaller firms, whereas *Content Generation*, *Data Analytics*, and *Service & Fulfillment* account for larger shares among larger firms. For example, Operational Guidance represents 41.7% of requests among bottom-quartile firms versus 33.4% among top-quartile firms, while the share of

Content Generation rises from 27.8% to 42.5%. These patterns illustrate differences in how firms of different sizes use the same AI system.

A.2 Data Description

A.2.1 Variable Definitions and Balance Check

Appendix Table A.1 defines the variables used in the firm-level analysis, covering sales performance, business capabilities, operating inputs, and information-related outcomes.

Appendix Table A.2 assesses covariate balance in the matched sample. We estimate propensity scores separately for each adoption cohort using pre-treatment firm characteristics and operating outcomes, including Level-1 category indicators. We calculate treated and control means using the matching weights. The matched groups are closely balanced on observable pre-treatment characteristics: all reported absolute standardized mean differences are below 0.10.

A.2.2 Construction of Baseline Information-Friction Measures

Table A.3 defines the six predetermined category-level measures of baseline information frictions used in Section 6.2. All measures are constructed using pre-AI information and standardized before entering the regressions.

Appendix B A Stylized Model of AI Adoption

This appendix develops the model summarized in Section 2.2 and provides the comparative statics and proofs. Platform AI raises firm sales through two channels, lower task inputs and better information, and by more in categories where pre-AI information is less precise. The model does not sign the size gradient. A benchmark in which every adopter realizes the same increments from the platform's tools yields larger proportional gains for lower-capability firms (Section B.3, Proposition 1). When what firms realize from these tools rises steeply enough with their capability, the gradient can reverse (Lemma 3).

B.1 Firms, Information, and AI Adoption

A continuum of firms operates in product category m. Firm i has pre-AI business capability $k_i > 0$, which follows a continuous distribution $H_m(k)$ on $[\underline{k}_m, \overline{k}_m]$. Capability captures the firm's operating knowledge, proprietary data, managerial capacity, and ability to perform recurring business tasks.

A firm may operate without AI, use platform-provided AI, or develop and deploy AI internally. Let $j \in \{0, P, S\}$ index these states: 0 denotes no AI, P denotes platform provision, and S denotes self-provision. The empirical analysis observes adoption of the platform system. Self-provision serves as a theoretical benchmark illustrating how AI returns may depend on firms' proprietary resources and organizational capabilities. The empirical analysis focuses on adoption of the platform AI system rather than a comparison of provision modes.

Firm i is uncertain about a payoff-relevant state θ_i, such as consumer demand, product-market fit, customer preferences, or transaction risk. The prior in category m is

$$\theta_i \sim N\left(\mu_i, \tau_{0m}^{-1}\right), \qquad \tau_{0m} > 0. \tag{B.1}$$

Before using AI, the firm observes a private operating signal

$$s_i = \theta_i + \varepsilon_i, \qquad \varepsilon_i \sim N\left(0, \tau(k_i)^{-1}\right), \qquad \tau'(k) > 0. \tag{B.2}$$

The firm's total pre-AI posterior precision is therefore

$$a_m(k) \equiv \tau_{0m} + \tau(k). \tag{B.3}$$

A lower τ_{0m} means that information is less precise in the category. Firms with the same capability can therefore face different information frictions across categories.

AI supplied under mode $j \in \{P, S\}$ provides an additional signal

$$z_{ij} = \theta_i + u_{ij}, \qquad u_{ij} \sim N\left(0, b_j(k_i)^{-1}\right), \tag{B.4}$$

where the signal errors are independent of the state and the firm's private signal error. The function $b_j(k) > 0$ measures the precision of the AI signal. Because platform provision combines platform-wide information with a standardized model and workflow, the benchmark is

$$b_P(k) = \bar{b}_P > 0. \tag{B.5}$$

With self-provision, AI quality may depend on proprietary data, technical personnel, and complementary investments. Thus, $b_S(k)$ may increase with k.

Suppose the firm chooses an action under quadratic decision loss. Posterior variance is $a_m(k)^{-1}$ before AI use and $[a_m(k) + b_j(k)]^{-1}$ after the firm observes the AI signal. Let $\gamma > 0$ convert the reduction in expected decision loss into effective business capability. The resulting information-mediated capability increment is

$$I_{jm}(k) = \gamma \left[\frac{1}{a_m(k)} - \frac{1}{a_m(k) + b_j(k)}\right] = \gamma \frac{b_j(k)}{a_m(k)\left[a_m(k) + b_j(k)\right]}. \tag{B.6}$$

Differentiating with respect to firm capability gives

$$I'_{jm}(k) = \frac{\gamma}{\left[a_m(k) + b_j(k)\right]^2}\left[b'_j(k) - \frac{b_j(k)\left[2a_m(k) + b_j(k)\right]a'_m(k)}{a_m(k)^2}\right]. \tag{B.7}$$

Equation (B.7) shows the central tradeoff. More precise pre-existing information reduces the added value of AI. At the same time, the AI signal may be more precise for higher-capability firms, either because they are better at using it or, under self-provision, because their proprietary data and complementary investments improve it. Under the platform benchmark, $b'_P(k) = 0$ and $a'_m(k) =$

$\tau'(k) > 0$, so

$$I'_{Pm}(k) < 0. \tag{B.8}$$

The value of platform-provided information also varies with the precision of the category information environment:

$$\frac{\partial I_{Pm}(k)}{\partial \tau_{0m}} = -\gamma \bar{b}_P \frac{2a_m(k) + \bar{b}_P}{a_m(k)^2 \left[a_m(k) + \bar{b}_P\right]^2} < 0. \tag{B.9}$$

Thus, the same platform signal improves decisions more in categories where pre-AI information is less precise.

AI may also add capability directly through content creation, demand discovery, product management, service, and analytics support. Let this increment be $D_j(k) > 0$. Effective market and decision capability is

$$x_{jm}(k) = k + D_j(k) + I_{jm}(k), \qquad x_{0m}(k) = k. \tag{B.10}$$

The two increments capture different channels. $D_j(k)$ represents capabilities embedded directly in the AI tool. $I_{jm}(k)$ represents the value of better information relative to what the firm already knows. For platform provision, the benchmark is

$$D_P(k) = \bar{D}_P > 0. \tag{B.11}$$

We impose $x'_{jm}(k) > 0$, so AI does not reverse firms' ranking by effective capability.

Let

$$\delta_{jm}(k) \equiv D_j(k) + I_{jm}(k) > 0. \tag{B.12}$$

Lemma 1 (Capability-return elasticity)**.** *The derivative of the proportional capability gain is*

$$\frac{\partial}{\partial k} \log \frac{x_{jm}(k)}{k} = \frac{k\delta'_{jm}(k) - \delta_{jm}(k)}{k\left[k + \delta_{jm}(k)\right]}. \tag{B.13}$$

Hence the proportional capability gain decreases with initial capability if and only if

$$\frac{k\delta'_{jm}(k)}{\delta_{jm}(k)} < 1. \tag{B.14}$$

Proof. Since $x_{jm}(k) = k + \delta_{jm}(k)$,

$$\frac{\partial}{\partial k} \log \frac{x_{jm}(k)}{k} = \frac{1 + \delta'_{jm}(k)}{k + \delta_{jm}(k)} - \frac{1}{k},$$

which simplifies to equation (B.13). Because the denominator is positive, the derivative is negative if and only if $k\delta'_{jm}(k) < \delta_{jm}(k)$, yielding equation (B.14). ∎

Condition (B.14) determines the sign of the capability gradient. It allows $\delta_{jm}(k)$ to rise with initial capability, as long as it rises less than proportionally. Under the platform benchmark, $D'_P(k) = 0$ and $I'_{Pm}(k) < 0$, so the condition holds. It fails if what firms realize from AI rises sufficiently with their capability: under platform provision, through implementation (Lemma 3); under self-provision, also through model quality and customization, which depend on proprietary data and organizational capability.

B.2 Operating Efficiency and Marginal Cost

Let $O_j(k) > 0$ denote the operating capability added under mode j, and define

$$r_j(k) = k + O_j(k), \qquad r_0(k) = k. \tag{B.15}$$

The task input required per unit of activity is

$$\ell_j(k) = \psi_j r_j(k)^{-\nu}, \qquad \nu > 0, \tag{B.16}$$

where ψ_j captures direct task automation. Without AI, this parameter is $\psi_0 > 0$. AI reduces the inputs required for a task, so

$$0 < \psi_j < \psi_0, \qquad j \in \{P, S\}. \tag{B.17}$$

Platform provision supplies operating modules through a standardized workflow. We use

$$O_P(k) = \bar{O}_P > 0 \tag{B.18}$$

as the benchmark. Self-provided operating capability may instead depend strongly on technical personnel, internal systems, and process redesign.

Lemma 2 (Operating-return elasticity)**.** *For any AI mode j with ψ_j independent of k,*

$$\frac{\partial}{\partial k} \log \frac{\ell_0(k)}{\ell_j(k)} = \nu \frac{kO_j'(k) - O_j(k)}{k\left[k + O_j(k)\right]}. \tag{B.19}$$

Hence proportional task-input savings decrease with initial capability if and only if

$$\frac{kO_j'(k)}{O_j(k)} < 1. \tag{B.20}$$

Proof. Using equation (B.16),

$$\log \frac{\ell_0(k)}{\ell_j(k)} = \log \frac{\psi_0}{\psi_j} + \nu \left[\log\left(k + O_j(k)\right) - \log k\right].$$

Differentiating gives equation (B.19); its sign is negative if and only if equation (B.20) holds. ∎

For the platform benchmark,

$$\log \frac{\ell_0(k)}{\ell_P(k)} = \log \frac{\psi_0}{\psi_P} + \nu \log \frac{k + \bar{O}_P}{k}, \tag{B.21}$$

which is positive and decreasing in k.

Marginal cost is

$$c_j(k; w_m) = \bar{c} + w_m \ell_j(k) = \bar{c} + w_m \psi_j r_j(k)^{-\nu}, \qquad \bar{c} > 0,\ w_m > 0, \tag{B.22}$$

where w_m is the price of task inputs and $\bar{c}$ captures marginal costs that AI does not affect.

Under the platform benchmark, define

$$a_0(k) \equiv \psi_0 k^{-\nu}, \qquad a_P(k) \equiv \psi_P (k + \bar{O}_P)^{-\nu}. \tag{B.23}$$

Since $a_0(k) > a_P(k)$,

$$\frac{\partial}{\partial w_m} \log \frac{c_0(k; w_m)}{c_P(k; w_m)} = \frac{\bar{c}\,[a_0(k) - a_P(k)]}{[\bar{c} + w_m a_0(k)]\,[\bar{c} + w_m a_P(k)]} > 0. \tag{B.24}$$

Task savings are therefore more valuable where labor and other inputs are more expensive.

To sign the capability gradient of marginal-cost savings, define

$$A(r, \psi) \equiv \frac{w_m \psi r^{-\nu-1}}{\bar{c} + w_m \psi r^{-\nu}}. \tag{B.25}$$

The function $A(r, \psi)$ is decreasing in r and increasing in ψ. Because $k + \bar{O}_P > k$ and $\psi_P < \psi_0$,

$$A(k + \bar{O}_P, \psi_P) < A(k, \psi_0). \tag{B.26}$$

It follows that

$$\frac{\partial}{\partial k} \log \frac{c_0(k; w_m)}{c_P(k; w_m)} = -\nu A(k, \psi_0) + \nu A(k + \bar{O}_P, \psi_P) < 0. \tag{B.27}$$

B.3 Demand, Pricing, and Firm-Level Returns

Within category m, firms sell differentiated varieties under CES demand. A firm's effective market capability shifts demand for its variety:

$$q_{ij} = Q_m x_{jm}(k_i)^{\eta_m} \left(\frac{p_{ij}}{\mathcal{P}_m}\right)^{-\sigma}, \qquad \sigma > 1,\ \eta_m > 0. \tag{B.28}$$

An atomistic firm takes the category aggregates Q_m and $\mathcal{P}_m$ as given. Constant-elasticity demand yields the price

$$p_{ij} = \mu c_j(k_i; w_m), \qquad \mu \equiv \frac{\sigma}{\sigma - 1}. \tag{B.29}$$

Firm sales are

$$S_j(k) = B_m x_{jm}(k)^{\eta_m} c_j(k; w_m)^{1-\sigma}, \qquad B_m \equiv Q_m \mathcal{P}_m^{\sigma} \mu^{1-\sigma}. \tag{B.30}$$

Relative to operating without AI, mode j generates the proportional sales return

$$\begin{aligned} R_j(k) &\equiv \log \frac{S_j(k)}{S_0(k)} \\ &= \eta_m \log \frac{x_{jm}(k)}{k} + (\sigma - 1) \log \frac{c_0(k; w_m)}{c_j(k; w_m)}. \end{aligned} \tag{B.31}$$

The first term captures the capability and information channel. The second captures the operating-cost channel.

Proposition 1 (Returns under the platform benchmark)**.** *Suppose platform provision satisfies equations* (B.5)*,* (B.11)*, and* (B.18)*, with* $x'_{Pm}(k) > 0$ *and* $0 < \psi_P < \psi_0$*. Then platform-provided AI:*

1. *increases market capability and reduces task input requirements and marginal cost;*
2. *raises firm sales,* $R_P(k) > 0$*; and*
3. *generates larger proportional capability, task-input, marginal-cost, and sales gains for firms*

with lower initial capability:

$$\frac{\partial}{\partial k}\log\frac{x_{Pm}(k)}{k} < 0, \tag{B.32}$$

$$\frac{\partial}{\partial k}\log\frac{\ell_0(k)}{\ell_P(k)} < 0, \tag{B.33}$$

$$\frac{\partial}{\partial k}\log\frac{c_0(k;w_m)}{c_P(k;w_m)} < 0, \tag{B.34}$$

$$R'_P(k) < 0. \tag{B.35}$$

Proof. Under the platform benchmark, $D'_P(k) = 0$ and $I'_{Pm}(k) < 0$. Hence $\delta'_{Pm}(k) < 0$, and Lemma 1 implies equation (B.32). Lemma 2, together with $O'_P(k) = 0$, implies equation (B.33). Equation (B.27) establishes equation (B.34). Platform provision raises $x_{Pm}(k)$ and lowers $c_P(k;w_m)$, so both terms in equation (B.31) are positive. Both terms are also decreasing in k, which proves $R_P(k) > 0$ and equation (B.35). ∎

Proposition 1 distinguishes proportional returns from absolute gains. Since

$$S_P(k) = S_0(k)e^{R_P(k)}, \tag{B.36}$$

the gross profit gain from platform AI is

$$G_P(k) = \frac{S_0(k)}{\sigma}\left[e^{R_P(k)} - 1\right]. \tag{B.37}$$

Differentiating shows that $G'_P(k) > 0$ if and only if

$$\frac{S'_0(k)}{S_0(k)} > -\frac{e^{R_P(k)}}{e^{R_P(k)} - 1}R'_P(k). \tag{B.38}$$

A smaller percentage gain can produce a larger absolute gain when it applies to a sufficiently large baseline. Thus, larger firms may be more likely to adopt even when smaller adopters receive larger proportional gains.

Proposition 1 describes one force behind the size gradient, duplication: standardized tools partly duplicate what high-capability firms already have and fill larger gaps for lower-capability firms. The benchmark gradient need not hold, because a second force, complementarity, can work in the opposite direction.

Implementation heterogeneity. Equations (B.5), (B.11), and (B.18) treat what the platform delivers as identical across adopters. What a firm realizes from these tools may still depend on its capability. Firms with stronger managers and organizations may be better at selecting which AI outputs to use, adapting them to their products, and acting on AI recommendations (Otis et al., 2026). To capture this, scale the realized platform increments by an implementation factor $\phi(k) > 0$ with $\phi'(k) \geq 0$:

$$D_P(k) = \phi(k)\,\bar{D}_P, \qquad O_P(k) = \phi(k)\,\bar{O}_P, \qquad b_P(k) = \phi(k)\,\bar{b}_P.$$

The benchmark in Proposition 1 is the case $\phi'(k) = 0$. Let $\varepsilon_\phi(k) \equiv k\,\phi'(k)/\phi(k)$ denote the elasticity of implementation with respect to initial capability.

Lemma 3 (Implementation heterogeneity)**.** *Suppose the platform increments are scaled by $\phi(k)$ as above, with $x'_{Pm}(k) > 0$ and $0 < \psi_P < \psi_0$.*

1. *If $\varepsilon_\phi(k) < 1$, all conclusions of Proposition 1 hold, including equations* (B.32)–(B.35).
2. *If $\varepsilon_\phi(k) > 1$, proportional task-input savings increase with k. If, in addition, $\bar{c} = 0$, proportional marginal-cost savings also increase with k.*
3. *If $\varepsilon_\phi(k)$ is sufficiently large, proportional capability gains increase with k. If both proportional capability gains and proportional marginal-cost savings increase with k, then $R'_P(k) > 0$.*

Proof. Task inputs. By construction, $kO'_P(k)/O_P(k) = \varepsilon_\phi(k)$. By Lemma 2, proportional task-input savings decrease with k if and only if $\varepsilon_\phi(k) < 1$.

Capability. Taking logs of equation (B.6) with $b_P(k) = \phi(k)\bar{b}_P$ and differentiating,

$$\frac{k\, I'_{Pm}(k)}{I_{Pm}(k)} = \varepsilon_\phi(k)\, \frac{a_m(k)}{a_m(k) + b_P(k)} - \frac{k\, a'_m(k)}{a_m(k)} - \frac{k\, a'_m(k)}{a_m(k) + b_P(k)} < \varepsilon_\phi(k),$$

since $a'_m(k) = \tau'(k) > 0$. Because $kD'_P(k)/D_P(k) = \varepsilon_\phi(k)$,

$$\frac{k\, \delta'_{Pm}(k)}{\delta_{Pm}(k)} = \frac{D_P(k)}{\delta_{Pm}(k)}\, \varepsilon_\phi(k) + \frac{I_{Pm}(k)}{\delta_{Pm}(k)}\, \frac{k\, I'_{Pm}(k)}{I_{Pm}(k)} < \varepsilon_\phi(k).$$

Hence condition (B.14) holds whenever $\varepsilon_\phi(k) < 1$, and Lemma 1 implies equation (B.32). Holding $\phi(k)$ fixed, the right-hand side of the last expression is linear and strictly increasing in $\phi'(k)$, so condition (B.14) fails when $\varepsilon_\phi(k)$ is sufficiently large.

Marginal cost. Let $r_P(k) = k + \phi(k)\bar{O}_P$. Differentiating equation (B.22),

$$\frac{\partial}{\partial k} \log \frac{c_0(k; w_m)}{c_P(k; w_m)} = -\nu A(k, \psi_0) + \nu A\big(r_P(k), \psi_P\big)\, r'_P(k).$$

Write $A(r, \psi) = s(r, \psi)/r$, where

$$s(r, \psi) \equiv \frac{w_m \psi r^{-\nu}}{\bar{c} + w_m \psi r^{-\nu}}$$

is the task-input share of marginal cost. The share s is decreasing in r and increasing in ψ, so $s(r_P(k), \psi_P) < s(k, \psi_0)$. Moreover,

$$\frac{r'_P(k)}{r_P(k)} < \frac{1}{k} \iff k + k\phi'(k)\bar{O}_P < k + \phi(k)\bar{O}_P \iff \varepsilon_\phi(k) < 1.$$

If $\varepsilon_\phi(k) < 1$, then $A(r_P, \psi_P)\, r'_P = s(r_P, \psi_P)\, r'_P/r_P < s(k, \psi_0)/k = A(k, \psi_0)$, which establishes equation (B.34). If $\bar{c} = 0$, then $s \equiv 1$, and the derivative has the sign of $r'_P/r_P - 1/k$, which is positive if and only if $\varepsilon_\phi(k) > 1$.

Sales. By equation (B.31), $R_P(k)$ is the sum of the proportional capability term and the proportional marginal-cost term. If $\varepsilon_\phi(k) < 1$, both terms are positive and decreasing in k, which

establishes $R_P(k) > 0$ and equation (B.35). If both terms are increasing in k, then $R'_P(k) > 0$. ∎

Lemma 3 shows that the sign of the size gradient depends on which force dominates. Duplication dominates when realized increments rise less than proportionally with capability, and complementarity dominates when they rise sufficiently more than proportionally. The model does not determine which case applies, so we treat the direction of the size gradient as an empirical question (Section 5). The cross-category comparative static does not depend on this. For any $b_j(k) > 0$,

$$\frac{\partial I_{jm}(k)}{\partial \tau_{0m}} = -\gamma\, b_j(k)\, \frac{2a_m(k) + b_j(k)}{a_m(k)^2\left[a_m(k) + b_j(k)\right]^2} < 0,$$

so the cross-market gradient in Prediction 3 holds for any $\phi(k)$ and under either mode of provision.

Self-provision is the case in which complementarity is likely to be strongest, because AI quality itself depends on firm-owned inputs. By Lemma 1, self-provided AI generates proportional market-capability returns that rise with k whenever

$$\frac{k\left[D'_S(k) + I'_{Sm}(k)\right]}{D_S(k) + I_{Sm}(k)} > 1. \tag{B.39}$$

For example, if $D_S(k) + I_{Sm}(k) = dk^\rho$ with $\rho > 1$, then proportional market-capability returns increase with k. Similarly, Lemma 2 implies that proportional task-input savings increase with capability whenever

$$\frac{kO'_S(k)}{O_S(k)} > 1. \tag{B.40}$$

If $\bar{c} = 0$, this condition also makes proportional marginal-cost savings rise with k. More generally, strong complementarities between self-provided AI and proprietary data, technical personnel, or organizational capital can outweigh the duplication effect of better information.

B.4 Adoption and the Upstream Location of Scale Economies

Let $F_j > 0$ denote the recurring burden of adopting mode $j \in \{P, S\}$, including monetary payments, learning, and workflow adjustment. We treat these burdens as determined by exogenous provision

and pricing rules. Platform provision typically implies

$$F_P < F_S, \tag{B.41}$$

because the platform centralizes model development, computing, data infrastructure, and workflow integration and shares them across many downstream users. With self-provision, firms bear most of the corresponding development and integration costs themselves. Equation (B.41) is a benchmark, not a necessary condition: firms with strong technical and organizational capabilities may face a relatively low burden of self-provision.

Variable profit equals $S_j(k)/\sigma$. The gross profit gain from mode j is

$$G_j(k; B_m) = \frac{S_j(k) - S_0(k)}{\sigma}. \tag{B.42}$$

The firm chooses

$$j^*(k) \in \arg\max \{0,\ G_P(k; B_m) - F_P,\ G_S(k; B_m) - F_S\}\ . \tag{B.43}$$

The self-provision option matters especially for high-capability firms, whose internal assets may strongly complement self-provided AI.

Proposition 2 (Platform adoption cutoff)**.** *Suppose $G_P(k; B_m)$ is continuous and strictly increasing in k, with*

$$G_P(\underline{k}_m; B_m) < F_P < G_P(\overline{k}_m; B_m).$$

Then there exists a unique cutoff k_P^* such that platform AI generates a nonnegative net gain relative to no AI if and only if $k \geq k_P^*$. The cutoff satisfies

$$G_P(k_P^*; B_m) = F_P, \tag{B.44}$$

and

$$\frac{\partial k_P^*}{\partial F_P} = \frac{1}{G_{P,k}(k_P^*; B_m)} > 0. \tag{B.45}$$

Proof. Continuity, strict monotonicity, and the crossing condition imply existence and uniqueness. Equation (B.45) follows from the implicit-function theorem. ∎

A lower F_P extends adoption to firms with lower initial capability. This extensive-margin access effect differs from the intensive-margin comparison in Proposition 1 and Lemma 3: a lower F_P brings in lower-capability firms whatever the sign of the size gradient among adopters.

B.5 Category-Level Market Expansion

Let $j(k) \in \{0, P\}$ denote the provision state of a firm with capability k, and let M_m denote the measure of active varieties. The effective CES category price index is

$$\mathcal{P}_m = \left[M_m \int x_{j(k)m}(k)^{\eta_m} p_{j(k)}(k)^{1-\sigma} \, dH_m(k) \right]^{\frac{1}{1-\sigma}}. \tag{B.46}$$

The index combines posted prices with the contribution of market capability to demand. It therefore measures the effective price of the category composite, rather than a price index adjusted for an intrinsic firm characteristic.

For every platform adopter, $x_{Pm}(k) > k$ and $p_P(k) < p_0(k)$. Because $1 - \sigma < 0$, platform adoption raises the integrand in equation (B.46) and therefore lowers $\mathcal{P}_m$.

Suppose upper-tier category demand is

$$Q_m = A_m \mathcal{P}_m^{-\varepsilon_m}, \qquad \varepsilon_m > 0. \tag{B.47}$$

Category sales are

$$X_m = \mathcal{P}_m Q_m = A_m \mathcal{P}_m^{1-\varepsilon_m}. \tag{B.48}$$

Proposition 3 (Market expansion)**.** *If a positive measure of firms uses platform-provided AI, then*

the effective category price index falls and category quantity increases. If $\varepsilon_m > 1$, aggregate category sales also increase.

Proof. Platform use raises $x_{jm}(k)$ and lowers $p_j(k)$ for every adopter. The integral inside equation (B.46) therefore rises. Since $1/(1-\sigma) < 0$, $\mathcal{P}_m$ falls. Equation (B.47) then implies that Q_m rises. Finally, equation (B.48) implies that X_m rises when $\varepsilon_m > 1$. ∎

The information structure yields one more category-level implication.

Corollary 1 (Information frictions and market expansion)**.** *Hold fixed the adoption set, the distribution of firm capability, and other category primitives. A reduction in the baseline category information precision τ_{0m} increases $I_{Pm}(k)$ for every adopter, and generates a larger increase in category quantity. If $\varepsilon_m > 1$, it also generates a larger increase in category sales.*

Proof. Equation (B.9) implies that lower τ_{0m} raises the information-mediated capability increment for every adopter. This raises $x_{Pm}(k)$ and therefore increases the integrand in equation (B.46). Because the outer exponent is negative, the effective price index falls by more. The remaining results follow from equations (B.47) and (B.48). ∎

Appendix C Instrumental-Variable Implementation

Instrument Construction. The task-demand component, D_i, is constructed from platform records for January through September 2023, before the AI tools were introduced. It averages two measures: the frequency of product-listing modifications and human customer-service time per order. Both measures are standardized within categories before averaging. The provincial digital-readiness component, $R_{p(i)}$, averages three measures from 2023: computers per 100 employees, websites per 100 firms, and the share of firms engaged in e-commerce, each normalized to [0, 1] using its minimum and maximum across provinces. The time-varying component, $A_{-c(i),h(i),t}$, is the share of firms active in month t in other Level-2 categories within firm i's Level-1 parent category that had used AI by month t. Each firm is assigned to a unique primary level-2 category based on its pre-AI operations, and this assignment remains fixed throughout the analysis.

Institutional Motivation and Robustness. The platform promotes AI adoption through both platform-wide training and industry-specific merchant communities. Industry teams and training groups are broadly organized around level-1 product sectors, facilitating information diffusion across distinct level-2 markets. In addition to established merchant groups, the platform operates dedicated AI training groups that provide instruction exclusively on AI tools. Outside major promotional events, routine business support and traffic incentives are generally allocated based on firm- or product-level characteristics rather than product categories. These institutional features motivate the use of related-category adoption as a measure of AI diffusion within a shared industry ecosystem. Table C.1 additionally includes province-by-month fixed effects to absorb regional shocks and product-category-by-month fixed effects to absorb contemporaneous shocks within the firm's own product market. Tables C.2 and C.3 report the corresponding first-stage estimates for the operating-cost and information-friction heterogeneity specifications, respectively.

Appendix D Robustness Analyses

D.1 Dynamic effects on Business Capabilities

This section presents additional event-study results for the firm-level outcomes discussed in the main text. Figures D.1–D.3 report dynamic effects on business capabilities.

D.2 Alternative Functional Form

For five outcomes with substantial mass at zero, our baseline specifications use PPML. As a robustness check, we reestimate the corresponding DID and IV specifications using $\log(1+x)$ transformations of the dependent variables. Table D.1 reports the results. Panel A presents OLS-DID estimates, while Panel B reports IV-2SLS estimates. Because $\log(1+x)$ estimates can be sensitive to the units of measurement when outcomes contain zeros (Chen and Roth, 2024), we interpret these results as checks on the direction and statistical significance of the baseline findings.

D.3 Staggered Difference-in-Differences

To address staggered treatment timing, we apply the imputation estimator of Borusyak, Jaravel, and Spiess (2024). It allows treatment effects to vary across firms and over time and avoids the weighting problems that conventional two-way fixed-effects estimators can create under staggered adoption. We estimate the model on the deduplicated firm-month panel. Table D.2 reports estimates that are qualitatively similar to the baseline findings.

D.4 Synthetic Difference-in-Differences

Following Arkhangelsky et al. (2021), we use synthetic difference-in-differences (SDID) as an additional robustness check for monthly sales. We estimate SDID separately within cells defined by Level-1 product category, baseline firm-size quartile, and adoption cohort. We use Level-1 categories to maintain sufficiently large donor pools, with firm-size quartiles defined as in the baseline heterogeneity analysis and adoption cohorts spanning July 2024 through June 2025. Within each cell, the donor pool consists of never-adopters in the same product category and firm-size quartile. The same never-adopter may serve as a control across multiple adoption cohorts, with unit and time weights estimated separately for each cell.

We aggregate the cell-specific estimates within each firm-size quartile, weighting each estimate by the number of treated firms in the corresponding cell. This gives greater weight to cells with more treated firms rather than weighting all cells equally. Figure D.4 reports the aggregated estimates, which are qualitatively similar to the baseline findings.

D.5 Sample-Definition Robustness

D.5.1 Relaxing the Continuous-Activity Restriction

The baseline sample requires firms to record positive sales in every month of the sample period. We relax this restriction to include firms with zero-sales months, while retaining all other sample restrictions and the baseline estimation procedures. Table D.3 shows that the main findings remain

robust. AI adoption increases sales, orders, and firm page views, while the estimated increase in average order prices is economically modest. Refund rates do not increase significantly, and 24-hour pickup performance improves.

D.5.2 Effects of AI Functionalities

We examine the two AI functionalities separately by restricting the treatment groups to single-function adopters. *Content Generation Only* firms adopt Content Generation but never Business Operation, while *Business Operation Only* firms adopt Business Operation but never Content Generation. Control firms adopt neither functionality. To align adoption cohorts with the baseline analysis, we exclude Content-Generation-only firms that first adopted between April and June 2024. All other sample restrictions and estimation procedures follow the baseline analysis.

Appendix Table D.4 shows that both functionalities independently increase sales, orders, and firm page views, with smaller sales gains than the 51.6% baseline estimate for joint adopters. Content Generation increases refund rates and reduces firm-liability disputes, but has no significant effect on 24-hour pickup performance. Business Operation reduces firm-liability disputes and improves pickup performance, consistent with its broader role in operational guidance and fulfillment.

D.6 Existing-Product Sales

We reestimate the baseline DID specification using monthly sales excluding products introduced in the same month. Table D.5 shows that sales from previously introduced products increase significantly following AI adoption, with larger gains among smaller firms. These findings indicate that the sales effects are not driven solely by contemporaneous new-product introductions.

D.7 Competitive Spillovers

This section describes the implementation of the three competitive-spillover tests in Section 3.4. The historical counterfactual is reported in Figure III; category-level exit and competitor exposure

are reported in Tables D.6 and D.7. All three provide little evidence that adopters gain at the expense of non-adopters.

Historical Counterfactual For the historical-counterfactual exercise, we calculate outcome-specific growth factors using aggregate outcomes in January–March 2024 relative to the same months of 2023, before the AI tools were introduced. We apply these factors to the corresponding 2023 outcomes for control firms, except for Product Negative Review Rate, which is left unadjusted because it is a rate rather than an activity-volume measure. Treated firms retain their original outcomes. We then re-estimate the baseline DID with firm-by-cohort and cohort-by-calendar-month fixed effects and category-clustered standard errors. Figure III reports the results.

Category Exit To assign each firm to a unique market, we define its primary operating category at each classification level as the category in which it generated the highest sales over the preceding year. We use transaction records through September 2025 to ensure a complete three-month exit window for every month in the analysis period.

We estimate exit rates under alternative market definitions. The Level-3 specifications include Level-3 category fixed effects and Level-2-category-by-month fixed effects, while the Level-2 specifications include Level-2 category fixed effects and Level-1-category-by-month fixed effects. Standard errors are clustered at the corresponding higher-level category. Table D.6 reports the estimates, which show no significant positive association between category-level AI adoption and exit rates, either overall or among non-adopters.

Competitor Exposure For each of the two Level-2 category-level exposure measures defined in the main text, we augment the baseline specification with the interaction

$$Treat_{ig} \times Post_{gt} \times E_{c(i)t},$$

together with the constituent lower-order terms. The regressions retain firm-by-cohort and cohort-by-calendar-month fixed effects and category-clustered standard errors. Table D.7 reports the estimates. The interaction coefficients are statistically insignificant for both exposure measures, providing little evidence that the treated-control sales gap widens with greater exposure to AI-enabled competitors.

Appendix E Heterogeneity in Capability Gains

This section provides additional evidence on how the effects of AI adoption vary with pre-treatment firm size. Table E.1 reports IV estimates for firm performance, while Table E.2 reports size-specific estimates using the staggered-DID imputation estimator. Figure E.1 examines heterogeneity in new-customer value.

Figure E.2 examines heterogeneity among firms adopting only one AI functionality, using the baseline firm-size cutoffs. The smallest firms experience the largest proportional sales gains from both Content Generation and Business Operation. The two groups also differ in composition: Content-Generation-only adopters tend to be larger, whereas Business-Operation-only adopters are more concentrated among smaller firms.

Figure E.3 examines heterogeneity when the continuous-activity restriction is relaxed. Firm-size quartiles are redefined within the expanded sample, with pre-treatment annual-sales cutoffs of RMB 13,126, RMB 30,259, and RMB 58,632. Sales gains remain substantial in the bottom three quartiles but are smaller and statistically insignificant in the top quartile. The extensive-margin effects, measured by the probability of positive monthly sales, follow a similar pattern. The expanded sample also contains substantially smaller firms: its 75th-percentile annual-sales cutoff (RMB 58,632) is below the baseline sample's median (RMB 63,088). Thus, at least three quarters of firms in the expanded sample would fall into the bottom two quartiles under the baseline classification. These findings extend the baseline size gradient to smaller firms with more intermittent sales activity.

Finally, Figures E.4–E.9 report heterogeneity in customer reach, product offerings, quality, operating inputs, information provision, and the role of baseline information frictions. Tables E.3–E.5 provide additional IV estimates for customer reach, product offerings, and quality.

Appendix F Cost-Equivalent Benchmarks

This appendix details the two cost-equivalent benchmarks reported in Table VII. Both benchmarks value activities supported by platform AI at conventional market prices, rather than measuring firms' actual expenditures or realized cost savings.

F.1 Content-Production Cost Benchmark

We use reference prices from a platform-administered survey conducted in August 2024 among firms operating for more than one year with sales exceeding RMB 2,000 in the preceding month. The survey collected more than 2,700 valid responses. We value conventional content production at RMB 70 per image and RMB 200 per promotional video, assuming that half of newly introduced products use a video.

Let N_i^0 denote firm i's number of new products and I_i^0 its average number of images per new product, both measured over the 12 months preceding the April 2024 launch. The baseline annual content-production cost equivalent is

$$C_{i,\text{content}}^0 = N_i^0(70I_i^0 + 0.5 \times 200). \tag{F.1}$$

The growth-adjusted benchmark incorporates the estimated increases in new-product introductions and images per new product following AI adoption:

$$C_{i,\text{content}}^1 = N_i^0(1 + g_{q(i)}^N)\left[70I_i^0(1 + g_{q(i)}^I) + 0.5 \times 200\right], \tag{F.2}$$

where $q(i)$ denotes firm i's pre-launch sales quartile, and g_q^N and g_q^I are the corresponding

estimated proportional increases. Thus, image-production costs reflect growth in both new products and images per product, while video-production costs adjust only for growth in new products.

F.2 Comparable Outsourced Operating Services

External e-commerce operation-service providers offer services overlapping with Business Operation, including data analysis, store management, marketing, and customer service. Table F.1 documents illustrative pricing models from service providers and listed companies.

Based on these examples, we use an annual fixed fee of RMB 15,000 and an assumed commission of 3% of sales. The baseline and growth-adjusted benchmarks are

$$C^0_{i,\text{service}} = 15{,}000 + 0.03G^0_i, \tag{F.3}$$

$$C^1_{i,\text{service}} = 15{,}000 + 0.03G^0_i(1 + g^G_{q(i)}), \tag{F.4}$$

where G^0_i is baseline annual GMV and g^G_q is the estimated proportional sales increase for size quartile q. The fixed fee is informed by advertised service prices, while the 3% commission is an assumed parameter. The qualitative size gradient is similar under alternative fee and commission assumptions.

Table A.1
Variable Definitions

Variable	Definition
Sales	Monthly gross merchandise value (GMV), in RMB.
Orders	Number of orders received by the firm during the month.
Average Order Price	Monthly sales divided by the number of orders, in RMB.
Firm Page Views	Number of views of all pages, including firm's storefront and product pages, during the month.
New Visitors	Number of unique consumers whose first-ever visit to the firm occurs during the month. First-ever status is determined using the platform's full historical records and is not restricted to the analysis period.
New Customers	Number of unique consumers whose first-ever purchase from the firm occurs during the month. First-ever status is determined using the platform's full historical records and is not restricted to the sample period.
New-Customer Sales	Monthly sales generated by new customers.
Number of New Categories	Number of the platform's Level-4 categories entered by the firm for the first time during the month.
Number of New Products	Number of products newly introduced by the firm during the month.
New-Product Sales	Monthly sales generated by products introduced during the month.
Number of Images	Number of product images newly uploaded by the firm during the month, including images for new products and new images uploaded for existing products. Both manually produced and AI-assisted images are included.
Images per New Product	Number of images newly uploaded for products introduced during the month, divided by the number of newly introduced products.
Images per Existing Product	Number of new or revised images uploaded for existing products during the month, divided by the number of existing active products.
Modifications per Product	Number of product-listing modifications (e.g., to product titles, descriptions, or attributes) divided by the number of active products.
Human Customer-Service Time per Order	Total human customer-service time during the month divided by the number of orders, in seconds.
Product-Title Length	Average number of characters in the titles of newly introduced products or products with modified titles during the month. Missing for firm-months with no such products.
Delivery Time	Platform-recorded average logistics delivery duration, in hours, for orders placed with the firm during the month.
Refund Rate	Number of orders placed during the month that are subsequently refunded, divided by the total number of orders placed during that month.
Firm-Reason Refund Rate	Number of orders placed during the month that are subsequently refunded for firm-attributed reasons, divided by the total number of orders placed during that month.
Product Negative Review Rate	Number of negative product reviews attributed to orders placed during the month, divided by the number of orders placed in that month.
Firm-Liability Dispute Rate	Number of platform-adjudicated disputes resulting in firm liability, divided by the number of orders placed in that month.
New-Product Success Rate	Share of products introduced during the month that generate at least one sale.
New-Customer Conversion Rate	Share of consumers making their first-ever visit to the firm during the month who also make their first-ever purchase from the firm within the same month.
24-Hour Pickup On-Time Rate	Share of orders picked up by the carrier within 24 hours of purchase.
New-Product Attribute Coverage	Average number of filled attributes for newly introduced products, divided by the maximum observed number of attributes in the same Level-2 category. Missing for firm-months without new products.

Notes: Refunds, reviews, and disputes are tracked through July 2026 and attributed to their corresponding order months.

Table A.2
Covariate Balance After Matching

Baseline Variable	Treated Mean	Control Mean	SMD
log(1 + Sales)	8.42	8.46	0.026
log(1 + Orders)	3.96	3.85	0.082
log(1 + Number of Active Products)	2.35	2.36	0.018
Tenure	6.71	7.08	0.080
Flagship-Store Type	0.044	0.045	0.005

Notes: This table reports covariate balance between treated firms and their matched controls. Propensity scores are estimated separately for each monthly adoption cohort using pre-treatment sales, orders, active products, tenure, firm type, and Level-1 product-category indicators; each treated firm is matched to the nearest never-adopter with replacement. Time-varying covariates are averaged over the cohort-specific pre-treatment period, while time-invariant characteristics are measured at baseline. Standardized mean differences (SMDs) are reported in absolute values. Flagship-Store Type equals one for flagship-store firms and zero for independent third-party firms.

Table A.3
Construction of Baseline Information-Friction Measures

Measure	Construction	Greater Friction
Conversion Rate	The number of orders divided by the number of firm views within third-level category c during the pre-AI baseline period.	Lower
Refund Rate	The number of refunded orders divided by the total number of orders within third-level category c during the pre-AI baseline period.	Higher
Firm Exit Rate	The number of firms with positive sales in month $t-1$ but no sales in months t, $t+1$, and $t+2$, divided by the number of firms with positive sales in month $t-1$, aggregated at the Level-3 category level. The latest exit month is January 2024, allowing the full three-month exit window to be observed by March 2024.	Higher
Sales Volatility	$$\text{Sales Volatility}_c = SD_y\left(\Delta \log \text{Sales}_{cy}\right), \qquad y = 2020, \ldots, 2023.$$ Following Comin and Philippon (2005), sales volatility is measured as the standard deviation of annual sales growth within Level-3 category c. Annual sales from 2019–2023 are used to construct the growth-rate series.	Higher
Product Survival Duration	$$\text{Product Survival Duration}_c = \frac{1}{N_c} \sum_{j \in \mathcal{J}_c} \left(T_j^{\text{last sale}} - T_j^{\text{listing}}\right).$$ For products first listed during the year preceding AI introduction and confirmed delisted by March 2024, survival duration is the number of months between initial listing and the last recorded sale. The category-level measure averages these completed durations, excluding products still active at the baseline cutoff.	Lower
Product Turnover Rate	The number of product entries and exits relative to the lagged stock of active products within Level-3 category c. The category-level baseline measure is constructed over the pre-AI period.	Higher

Notes: This table defines the six predetermined category-level measures used to proxy for information frictions. All measures are constructed at the Level-3 category level before the introduction of platform AI. Conversion Rate, Refund Rate, Firm Exit Rate, Product Survival Duration, and Product Turnover Rate use data from the year preceding AI introduction; Sales Volatility is the standard deviation of annual category sales growth over 2020-2023, constructed from sales data beginning in 2019. All measures use information available through March 2024, before the platform-wide launch of the AI system. Lower Conversion Rate and shorter Product Survival Duration indicate greater information frictions, whereas higher values of Refund Rate, Firm Exit Rate, Sales Volatility, and Product Turnover Rate indicate greater frictions. All six measures are standardized before entering the regressions.

Table C.1
IV Robustness: Time-Varying Controls for Baseline Firm Characteristics

	Sales		
	(1)	(2)	(3)
Adoption	0.634***	0.676***	0.674***
	(0.0257)	(0.0278)	(0.0279)
Firm FE	✓	✓	✓
Province × Month FE	✓	✓	✓
Level-2 Category × Month FE	✓	✓	✓
Baseline Sales × Month		✓	✓
Baseline Active Products × Month			✓
Baseline Firm Tenure × Month			✓
First-stage KP F-statistic	444.17	441.22	463.14
Observations	4,974,574	4,974,574	4,974,574

Notes: This table reports instrumental-variable estimates of the effect of AI adoption on monthly sales, transformed using $\log(1+x)$. AI Adoption is instrumented using the exposure-based instrument, which interacts pre-AI task demand, provincial digital readiness, and leave-Level-2-category-out AI adoption in related product categories. All specifications include firm, province-by-month, and Level-2-product-category-by-month fixed effects. Column (2) additionally interacts baseline annual sales with calendar-month indicators. Column (3) further interacts baseline active products and firm tenure with calendar-month indicators. All three firm characteristics are measured before the introduction of AI, with sales and active products transformed using $\log(1+x)$. These interactions allow firms with different pre-AI characteristics to follow different calendar-time paths. Standard errors, clustered at the Level-2 product-category level, are reported in parentheses. The bottom rows report the Kleibergen–Paap rk Wald first-stage F-statistic and the number of observations. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C.2
First-Stage Estimates for Operating-Cost Specifications

Baseline Characteristic	(1) Content Input	(2) Service Input	(3) Combined Input	(4) Local Wage Level
Panel A. Dependent Variable: AI Adoption				
Adoption IV	2.276***	2.254***	2.267***	3.451***
	(0.0756)	(0.0753)	(0.0750)	(0.109)
Interaction IV	−0.157***	−0.0505	−0.230***	−0.885***
	(0.0186)	(0.0317)	(0.0208)	(0.0271)
Panel B. Dependent Variable: AI Adoption × Baseline Characteristic				
Adoption IV	−0.0484	0.0725***	−0.125***	−0.621***
	(0.0344)	(0.0203)	(0.0337)	(0.0645)
Interaction IV	1.820***	5.690***	3.025***	4.135***
	(0.289)	(0.306)	(0.245)	(0.0722)
KP F-statistic	459.35	447.93	469.53	506.71
Min SW F-statistic	393.66	697.58	551.95	1,283.01

Notes: This table reports the first-stage regressions for the IV heterogeneity specifications in Table VIII. Each column corresponds to a different standardized baseline operating-cost measure, as indicated by the column heading. Panel A uses AI Adoption, an indicator equal to one in and after a firm's adoption month, as the dependent variable. Panel B uses the interaction between AI Adoption and the corresponding baseline characteristic as the dependent variable. Both first-stage equations include the Adoption IV and its interaction with the baseline characteristic (Interaction IV) as excluded instruments. All specifications include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. The bottom rows report the Kleibergen–Paap rk Wald F-statistic and the minimum Sanderson–Windmeijer conditional F-statistic across the two endogenous regressors. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C.3
First-Stage Estimates for Information-Friction Heterogeneity Specifications

Baseline Characteristic	(1) Conversion Rate	(2) Refund Rate	(3) Firm Exit Rate	(4) Sales Volatility	(5) Product Survival Duration	(6) Product Turnover Rate
Panel A. Dependent variable: AI Adoption						
Adoption IV	2.123***	2.891***	2.274***	2.237***	2.883***	2.241***
	(0.104)	(0.0966)	(0.0651)	(0.153)	(0.0667)	(0.152)
Interaction IV	−2.091**	−0.0551	−0.106	0.175***	0.0324	0.129**
	(0.958)	(0.0698)	(0.118)	(0.0205)	(0.0432)	(0.0504)
Panel B. Dependent variable: AI Adoption × Baseline Characteristic						
Adoption IV	−0.471**	−0.0713	−0.402***	−0.0101	0.0322	0.147
	(0.222)	(0.0617)	(0.137)	(0.0725)	(0.0444)	(0.143)
Interaction IV	6.157***	2.869***	3.938***	5.342***	2.913***	4.910***
	(0.345)	(0.126)	(0.270)	(0.226)	(0.0814)	(0.320)
Firm FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
KP F-statistic	126.84	403.13	508.05	104.78	684.26	105.96
Min SW F-statistic	196.74	864.87	530.15	211.10	1,285.33	231.13
Observations	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574

Notes: This table reports the first-stage regressions for the IV heterogeneity specifications in Table X. Each column corresponds to a different standardized baseline information-friction measure, as indicated by the column heading. Panel A uses AI Adoption, an indicator equal to one in and after a firm's adoption month, as the dependent variable. Panel B uses the interaction between AI Adoption and the corresponding baseline characteristic as the dependent variable. Both first-stage equations include the Adoption IV and its interaction with the baseline characteristic (Interaction IV) as excluded instruments. All specifications include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. The bottom rows report the Kleibergen–Paap rk Wald F-statistic and the minimum Sanderson–Windmeijer conditional F-statistic across the two endogenous regressors. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table D.1
Alternative Functional Form: Log(1+x) Estimates

Variables	(1) Number of New Products	(2) Number of New Categories	(3) New-Product Sales	(4) Number of Images	(5) Modifications per Product
Panel A. OLS-DID estimates					
Treat × *Post*	0.149***	0.0243***	0.497***	0.336***	0.210***
	(0.0103)	(0.00158)	(0.0264)	(0.0125)	(0.00917)
Firm-Cohort FE	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓
Observations	6,020,256	6,020,256	6,020,256	6,020,256	6,020,256
R-squared	0.507	0.203	0.462	0.443	0.375
Panel B. IV-2SLS estimates					
AI Adoption	0.480***	0.177***	1.143***	0.763***	0.219***
	(0.0259)	(0.0113)	(0.0762)	(0.0417)	(0.0371)
Firm FE	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
First-stage	2.255***	2.255***	2.255***	2.255***	2.255***
KP F-statistic	207.80	207.80	207.80	207.80	207.80
Observations	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574

Notes: This table reports alternative estimates of the effects of platform AI adoption on product offerings and operating activities. All dependent variables are transformed using $\log(1 + x)$. Panel A reports OLS estimates using the baseline matched-cohort difference-in-differences design, with firm-by-cohort and cohort-by-calendar-month fixed effects. Treat × Post equals one for treated firms beginning in their adoption month. Panel B reports IV-2SLS estimates with firm and calendar-month fixed effects, instrumenting AI Adoption with the instrument in equation (3). Panel B also reports the first-stage coefficient and Kleibergen-Paap rk Wald F-statistic. Standard errors are clustered at the Level-2 product-category level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table D.2
Staggered DID: Firm Performance and Business Capabilities

	Firm Performance			Business Capabilities				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables	Sales	Orders	Average Order Price	Firm Page Views	Number of New Products	Product Negative Review Rate	Firm-Liability Dispute Rate	24-Hour Pickup On-Time Rate
AI Adoption	0.337***	0.317***	−0.000297	0.319***	0.0657***	−0.478***	−0.154***	0.0167***
	(0.00343)	(0.00315)	(0.00159)	(0.00350)	(0.00149)	(0.0143)	(0.0133)	(0.000603)
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574

Notes: This table reports robustness estimates using the imputation estimator of Borusyak, Jaravel, and Spiess (2024), with treatment beginning in each firm's first AI-adoption month. Sales, Orders, Average Order Price, Firm Page Views, and Number of New Products are transformed using $\log(1 + x)$. Product Negative Review Rate, Firm-Liability Dispute Rate, and 24-Hour Pickup On-Time Rate are normalized by their matched-control pre-adoption means, so their coefficients measure changes relative to the corresponding benchmark. The imputation model includes firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table D.3
Effects of AI Adoption Allowing Zero-Sales Months

Variables	(1) Sales	(2) Orders	(3) Average Order Price	(4) Firm Page Views	(5) Refund Rate	(6) Firm-Liability Dispute Rate	(7) 24-Hour Pickup On-Time Rate
Treat × *Post*	0.560***	0.351***	0.0172***	0.266***	0.0207	−0.254***	0.0274***
	(0.0574)	(0.0315)	(0.00499)	(0.0375)	(0.0242)	(0.0146)	(0.00148)
Firm-Cohort FE	✓	✓	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓	✓	✓
Observations	7,397,984	7,397,984	3,490,917	7,397,984	3,490,917	3,490,917	3,490,917
R-squared	0.300	0.331	0.790	0.299	0.350	0.255	0.506

Notes: This table reports matched difference-in-differences estimates after relaxing only the baseline requirement that firms have positive sales in every sample month. Treatment and control definitions, cohort construction, and matching otherwise follow the baseline analysis. Sales, Orders, Average Order Price, and Firm Page Views are transformed using $\log(1 + x)$. Refund Rate, Firm-Liability Dispute Rate, and 24-Hour Pickup On-Time Rate are normalized by their matched-control pre-adoption means, so $100\hat{\beta}$ gives the percentage change relative to the corresponding benchmark. Average Order Price and the three rate outcomes are defined only for firm-months with at least one order and therefore use the positive-order subsample. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table D.4
Effects of Single-Function AI Adoption

Variables	(1) Sales	(2) Orders	(3) Average Order Price	(4) Firm Page Views	(5) Refund Rate	(6) Firm-Liability Dispute Rate	(7) 24-Hour Pickup On-Time Rate
Panel A. Content Generation Only							
Treat × *Post*	0.225***	0.219***	−0.00477	0.209***	0.107***	−0.0567**	0.00104
	(0.0106)	(0.0125)	(0.00510)	(0.0142)	(0.0248)	(0.0261)	(0.00151)
Observations	3,983,034	3,983,034	3,983,034	3,983,034	3,983,034	3,983,034	3,983,034
R-squared	0.805	0.843	0.862	0.854	0.559	0.202	0.508
Panel B. Business Operation Only							
Treat × *Post*	0.205***	0.231***	−0.0451***	0.223***	0.0104	−0.148***	0.0196***
	(0.00867)	(0.0100)	(0.00800)	(0.00912)	(0.0101)	(0.0225)	(0.00183)
Observations	2,256,254	2,256,254	2,256,254	2,256,254	2,256,254	2,256,254	2,256,254
R-squared	0.691	0.739	0.825	0.797	0.335	0.211	0.486
Firm-by-Cohort FE	✓	✓	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓	✓	✓

Notes: This table reports matched difference-in-differences estimates of adopting a single AI functionality. Panel A includes firms that adopt Content Generation but never Business Operation; Panel B includes firms that adopt Business Operation but never Content Generation. Control firms adopt neither functionality. Treatment begins in the month of first adoption of the respective functionality. To align with the baseline cohort window, Content-Generation-only firms first adopting between April and June 2024 are excluded. Outcomes in columns (1)–(4) are transformed using $\log(1+x)$. Outcomes in columns (5)–(7) are normalized by their matched-control pre-adoption means, so $100\hat{\beta}$ gives the percentage change relative to those means. All regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Standard errors, clustered at the Level-2 product-category level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table D.5
Effects of AI Adoption on Existing-Product Sales

Variables	(1) All Firms	(2) Bottom	(3) Lower-Middle	(4) Upper-Middle	(5) Top
Panel A. Difference-in-differences estimates					
Treat × Post	0.359***	0.411***	0.238***	0.280***	0.0445
	(0.0125)	(0.0168)	(0.0159)	(0.0165)	(0.0269)
Firm-Cohort FE	✓	✓	✓	✓	✓
Cohort-by-Month FE	✓	✓	✓	✓	✓
% Change	43.2%	50.8%	26.9%	32.3%	–
Observations	6,020,256	1,505,086	1,505,042	1,505,042	1,505,086
R-squared	0.619	0.418	0.370	0.421	0.566
Panel B. Instrumental-variable estimates					
AI Adoption	0.787***	0.988***	0.803***	0.828***	0.604***
	(0.0284)	(0.0324)	(0.0282)	(0.0455)	(0.0470)
Firm FE	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
First-stage	2.26***	5.22***	4.97***	4.26***	4.32***
KP F-statistic	207.80	662.13	452.00	246.24	278.24
Observations	4,974,574	1,332,364	1,289,178	1,169,278	1,183,754

Notes: This table examines whether the sales gains from AI adoption extend beyond contemporaneous new-product introductions. Existing-Product Sales is defined as monthly firm sales excluding sales from products introduced in the same month. Products introduced in earlier months remain included, regardless of whether they were introduced before or after AI adoption. Column (1) reports estimates for the full matched sample, while columns (2)–(5) report estimates by quartile of pre-launch annual sales, using the baseline size cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. The dependent variable is transformed using $\log(1 + x)$. Panel A reports matched difference-in-differences estimates with firm-by-cohort and cohort-by-calendar-month fixed effects. The % Change row reports $100[\exp(\hat{\beta}) - 1]$ for statistically significant estimates. Panel B reports instrumental-variable estimates using the exposure-based instrument defined in equation (3), with firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 product-category level. The bottom rows of Panel B report the first-stage coefficient and Kleibergen-Paap rk Wald F-statistic. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table D.6
Category-Level AI Adoption and Firm Exit

Variables	All Firms		Non-Adopters	
	(1) Level-3 Category Exit Rate	(2) Level-2 Category Exit Rate	(3) Level-3 Category Exit Rate	(4) Level-2 Category Exit Rate
Category-Level AI Adoption Rate	0.0152	−0.0565	0.0174	0.00647
	(0.0114)	(0.0445)	(0.0113)	(0.0221)
Category FE	✓	✓	✓	✓
Level-2 Category × Month FE	✓		✓	
Level-1 Category × Month FE		✓		✓
R-squared	0.799	0.863	0.791	0.773

Notes: This table examines the relationship between AI diffusion within a product category and firm exit. Category-Level AI Adoption Rate is the number of active firms that have adopted AI by month t, divided by the total number of active firms in the category in that month. Exit in month t is defined as having positive sales in $t-1$ and no sales in t, $t+1$, or $t+2$. Columns (1) and (2) calculate the overall exit rate. Columns (3) and (4) calculate the non-adopter exit rate, restricting both the numerator and denominator to firms that had not adopted AI by month t. Columns (1) and (3) use Level-3 markets with Level-3 category and Level-2-category-by-month fixed effects; columns (2) and (4) use Level-2 markets with Level-2 category and Level-1-category-by-month fixed effects. Standard errors are clustered at the corresponding higher-level category. The number of category-month observations is not reported due to the platform's data-disclosure restrictions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table D.7
Firm-Level Sales Effects by Category-Level AI Adoption

Variables	(1) Sales	(2) Sales
Treat ×*Post*	0.416***	0.418***
	(0.0117)	(0.0127)
Treat ×*Post* × *Large-Firm AI Adoption Exposure*	0.0184	
	(0.0272)	
Treat ×*Post* × *Overall AI Adoption Rate*		0.0220
		(0.0170)
Firm-Cohort FE	✓	✓
Cohort-by-Month FE	✓	✓
Observations	6,020,256	6,020,256
R-squared	0.664	0.664

Notes: This table tests whether the firm-level effect of AI adoption varies with exposure to AI-enabled competitors in the firm's Level-2 category. The dependent variable is monthly sales, transformed using $\log(1+x)$. Column (1) interacts Treat × Post with Large-Firm AI Adoption Exposure, defined as cumulative AI adopters with annual sales above RMB 5 million divided by the current number of active firms in the category. Column (2) instead uses the cumulative adoption rate among all firms. Both exposure measures are standardized. All constituent lower-order terms of the interactions are included but omitted from the table. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table E.1
IV: Heterogeneity in Firm Performance

Variables	(1) Sales	(2) Orders	(3) Average Order Price
Adoption	1.256***	1.119***	0.0767***
	(0.0634)	(0.0564)	(0.0155)
Adoption × Lower-Middle	−0.442***	−0.368***	−0.0523***
	(0.0201)	(0.0197)	(0.00484)
Adoption × Upper-Middle	−0.664***	−0.545***	−0.0933***
	(0.0275)	(0.0256)	(0.00704)
Adoption × Top	−0.895***	−0.737***	−0.127***
	(0.0318)	(0.0322)	(0.00773)
Firm FE	✓	✓	✓
Month FE	✓	✓	✓
KP F-statistic	54.32	54.32	54.32
Observations	4,974,574	4,974,574	4,974,574

Notes: This table reports instrumental-variable estimates of the effects of AI adoption on sales performance by pre-treatment firm-size quartile. Sales, Orders, and Average Order Price are transformed using $\log(1 + x)$. The Bottom quartile is the omitted group; the coefficient on *Adoption* therefore gives the effect for bottom-quartile firms, while *Adoption* × Lower-Middle, Upper-Middle, and Top report differences relative to that group. The size cutoffs are RMB 24,460, RMB 63,088, and RMB 147,716 in pre-treatment annual sales. *Adoption* and its interactions with the three size indicators are instrumented by the baseline adoption instrument and its corresponding interactions. Regressions include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level. The table also reports the Kleibergen-Paap rk Wald F-statistic. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table E.2
Staggered DID: Heterogeneous Effects on Sales

	(1)	(2)	(3)	(4)
	Sales			
	Bottom	Lower-Middle	Upper-Middle	Top
AI Adoption	0.524***	0.369***	0.340***	0.220***
	(0.00680)	(0.00643)	(0.00673)	(0.00820)
Firm FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
Observations	1,332,364	1,289,178	1,169,278	1,183,754

Notes: This table reports estimates from the Borusyak, Jaravel, and Spiess (2024) imputation estimator separately by pre-treatment firm-size quartile. The dependent variable is monthly sales, transformed using $\log(1 + x)$, and treatment begins in each firm's first AI-adoption month. Columns (1)-(4) correspond to the Bottom, Lower-Middle, Upper-Middle, and Top quartiles, defined using annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. All specifications include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level and reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table E.3
IV: Heterogeneity in Customer Reach

Variables	(1) Firm Page Views	(2) New Visitors	(3) New Customers	(4) New-Customer Conversion Rate	(5) New-Customer Sales
Adoption	1.189^{***}	1.205^{***}	1.199^{***}	0.112^{***}	1.647^{***}
	(0.0732)	(0.0748)	(0.0572)	(0.0208)	(0.0830)
Adoption × Lower-Middle	-0.367^{***}	-0.360^{***}	-0.388^{***}	-0.0621^{***}	-0.554^{***}
	(0.0227)	(0.0231)	(0.0214)	(0.0106)	(0.0222)
Adoption × Upper-Middle	-0.544^{***}	-0.534^{***}	-0.595^{***}	-0.101^{***}	-0.849^{***}
	(0.0294)	(0.0288)	(0.0257)	(0.0100)	(0.0282)
Adoption × Top	-0.742^{***}	-0.733^{***}	-0.810^{***}	-0.0992^{***}	-1.126^{***}
	(0.0386)	(0.0369)	(0.0314)	(0.0139)	(0.0326)
Firm FE	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓
KP F-statistic	54.32	54.32	54.32	54.32	54.32
Observations	4,974,574	4,974,574	4,974,574	4,974,242	4,974,574

Notes: This table reports instrumental-variable estimates of the effects of AI adoption on customer reach by pre-treatment firm-size quartile. Firm Page Views, New Visitors, New Customers, and New-Customer Sales are transformed using $\log(1 + x)$. New-Customer Conversion Rate is the share of first-time visitors who make their first purchase from the firm within the same month, and is normalized by its matched-control pre-adoption mean. The Bottom quartile is omitted, so the *Adoption* coefficient gives the effect for bottom-quartile firms and the interaction coefficients give differences relative to that group. Size cutoffs are RMB 24,460, RMB 63,088, and RMB 147,716. *Adoption* and its size interactions are instrumented by the baseline adoption instrument and its corresponding interactions. Regressions include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level. The table also reports the Kleibergen-Paap rk Wald F-statistic.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table E.4
IV: Heterogeneity in Product Offerings

Variables	(1) Number of New Categories	(2) Number of New Products	(3) New-Product Success Rate	(4) New-Product Sales
Adoption	1.404***	0.862***	0.186***	2.530***
	(0.0608)	(0.162)	(0.0369)	(0.272)
Adoption × Lower-Middle	−0.323***	−0.139	−0.107***	−1.056***
	(0.0329)	(0.108)	(0.0120)	(0.196)
Adoption × Upper-Middle	−0.525***	−0.421***	−0.151***	−1.603***
	(0.0509)	(0.113)	(0.0132)	(0.237)
Adoption × Top	−0.610***	−0.563***	−0.184***	−1.958***
	(0.0433)	(0.119)	(0.0162)	(0.234)
Firm FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
Observations	3,531,814	4,350,324	1,509,819	4,251,456

Notes: This table reports instrumental-variable estimates of the effects of platform AI adoption on firms' product offerings by pre-treatment firm-size quartile. Columns (1), (2), and (4) are estimated using IV-PPML, while column (3) uses linear IV. New-Product Success Rate is defined as the share of newly introduced products generating at least one sale. It is observed only in firm-months with new products and is normalized by its matched-control pre-adoption mean. The Bottom quartile is omitted, so the *Adoption* coefficient gives the effect for bottom-quartile firms, and the interaction coefficients give differences relative to that group. Size cutoffs are RMB 24,460, RMB 63,088, and RMB 147,716. *Adoption* and its size interactions are instrumented using the baseline adoption instrument and its corresponding interactions. Regressions include firm and calendar-month fixed effects. IV-PPML estimation excludes separated observations and singleton fixed-effect groups where applicable. Standard errors are clustered at the Level-2 category level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table E.5
IV: Heterogeneity in Product and Service Quality

Variables	(1) Refund Rate	(2) Consumer-Reason Refund Rate	(3) Firm-Reason Refund Rate	(4) Product Negative Review Rate	(5) Firm-Liability Dispute Rate	(6) 24-Hour Pickup On-Time Rate	(7) Delivery Time
Adoption	0.222**	0.403***	−0.230***	−0.804*	−0.606***	0.0543***	−0.0296***
	(0.107)	(0.133)	(0.0545)	(0.468)	(0.129)	(0.0165)	(0.00614)
Adoption × Lower-Middle	0.00202	−0.0732***	0.00585	0.108	0.217***	−0.0151***	0.00705***
	(0.0168)	(0.0226)	(0.0233)	(0.0666)	(0.0371)	(0.00279)	(0.00193)
Adoption × Upper-Middle	0.00109	−0.0678**	−0.0143	0.263***	0.353***	−0.0157***	0.00656***
	(0.0225)	(0.0279)	(0.0243)	(0.0895)	(0.0469)	(0.00429)	(0.00240)
Adoption × Top	−0.0336	−0.0900**	0.0408	0.146	0.411***	−0.0315***	0.00458
	(0.0259)	(0.0350)	(0.0289)	(0.133)	(0.0534)	(0.00508)	(0.00312)
Firm FE	✓	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓	✓
KP F-statistic	54.32	54.32	54.32	54.32	54.32	54.32	54.32
Observations	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574	4,974,574

Notes: This table reports instrumental-variable estimates of the effects of AI adoption on product, service, and fulfillment quality by pre-treatment firm-size quartile. The rate outcomes in columns (1)-(6) are normalized by their matched-control pre-adoption means; Delivery Time in column (7) is transformed using $\log(1 + x)$. The Bottom quartile is omitted, so the *Adoption* coefficient gives the effect for bottom-quartile firms and the interaction coefficients give differences relative to that group. Size cutoffs are RMB 24,460, RMB 63,088, and RMB 147,716. *Adoption* and its size interactions are instrumented by the baseline adoption instrument and its corresponding interactions. Regressions include firm and calendar-month fixed effects. Standard errors are clustered at the Level-2 category level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table F.1
Examples of Pricing Models for E-Commerce Operation Services

Provider	Pricing Structure
Guangzhou Waimeng	Operating-service packages for a B2B e-commerce platform, priced at RMB 1,000, 1,500, 2,000, and 3,000 per month. Services range from basic store optimization to data and advertising management, product listing, and full-store operation.
Yijia Software	An enhanced annual operating-service package for a domestic B2B e-commerce platform, advertised at RMB 18,800 per year.
Ruoyuchen	A contractually agreed fixed fee and/or a negotiated proportion of sales after deducting agreed costs, used to determine monthly operating-service revenue.
Baozun	Fixed fees and/or variable fees based on merchandise sales, the number of orders fulfilled, or other contractual factors.

Notes: This table provides selected examples of pricing structures for external e-commerce operating services that overlap with activities supported by the platform AI system, including listing management, marketing, data analysis, customer service, and operating guidance. The examples are illustrative rather than exhaustive, and the services are not identical to the AI system studied in the paper. The advertised packages and disclosed pricing structures motivate the fixed-fee component and two-part structure of the benchmark in Table VII. The assumed 3% sales commission is a calibration parameter rather than an observed average commission across these providers. The benchmark does not represent firms' actual expenditures or realized savings from AI adoption.

Sources: Guangzhou Waimeng, service packages; Yijia Software, annual package; Ruoyuchen, 2025 Annual Report; Baozun, 2024 Form 20-F.

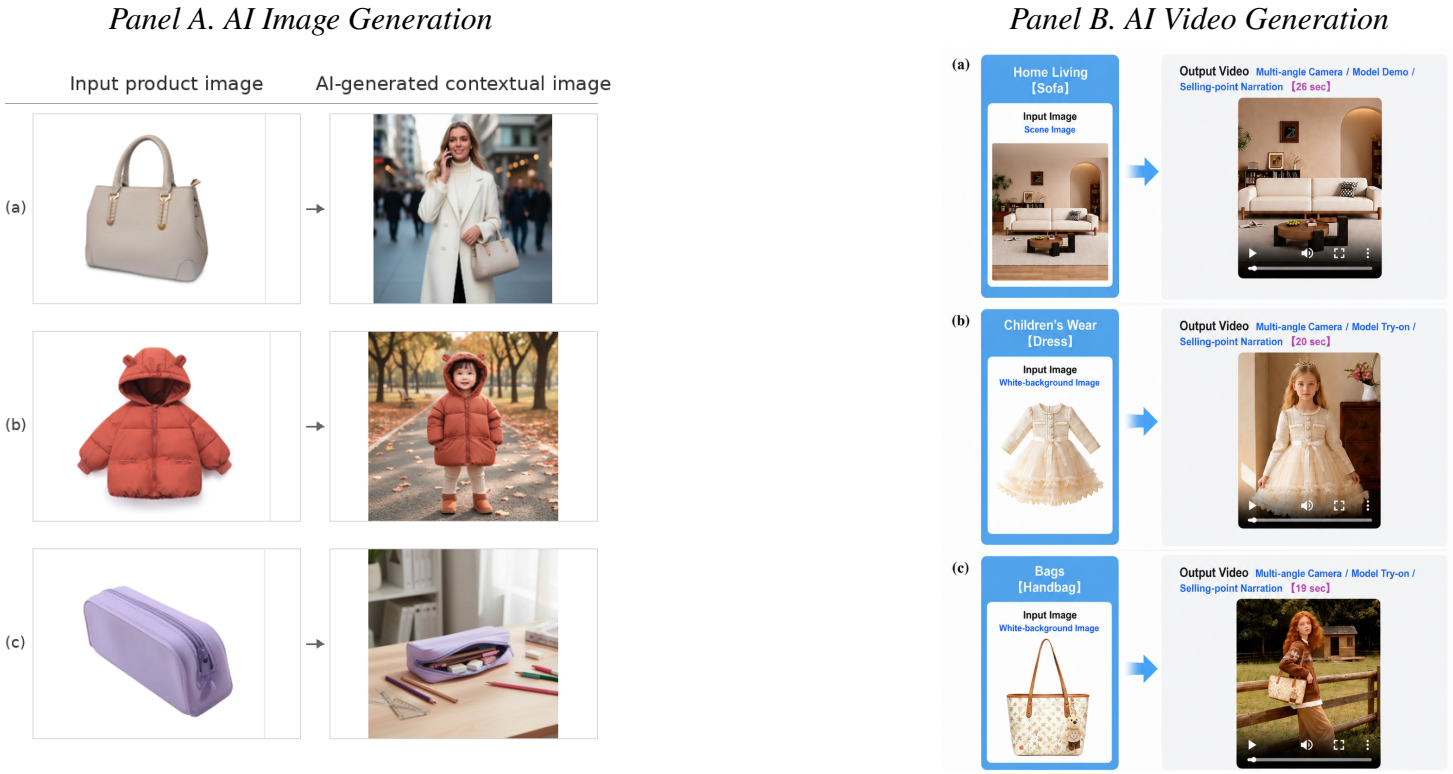


Figure A.1
Examples of AI-Assisted Content Generation

Notes: This figure illustrates two applications of the platform's Content Generation functionality. Panel A shows original product images (left) and corresponding AI-generated or edited listing images (right), including changes to backgrounds, models, and complementary objects. Panel B shows source product images (left) and examples of AI-generated promotional videos (right), which may include product demonstrations, usage scenes, camera movements, and narration. Both applications preserve the focal product. All examples are anonymized.



Figure A.2
Examples of Business-Operation Applications

Notes: This figure presents selected applications of the Business Operation functionality. Panels B-C illustrate Operational Guidance: firms can query platform rules and receive context-specific recommended actions. Panels D-F illustrate Data Analytics, including automated daily performance reports, market-trend analysis, and competitor analysis. Panels A, G, and I illustrate Market & Product Management, including product-title optimization, targeted customer outreach, and product or supplier discovery. Panel H illustrates Service & Fulfillment by monitoring logistics and after-sales exceptions that require firm attention. The applications combine firm-specific operating records with platform-wide information. All examples are anonymized.

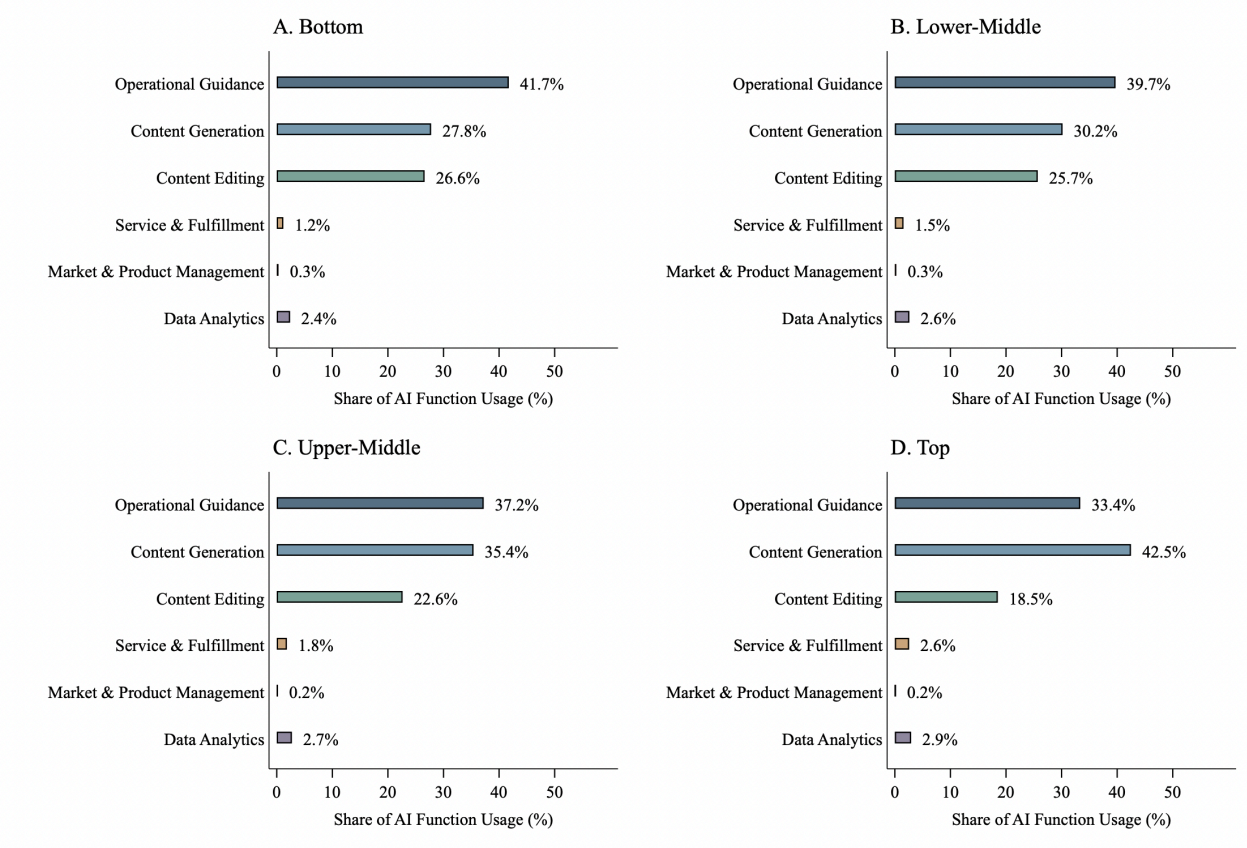


Figure A.3
AI Function Usage by Firm Size

Notes: This figure reports the composition of firm-initiated AI usage across the system's six business functions by pre-treatment firm-size quartile. Within each size group, each bar gives the share of recorded firm-initiated requests assigned to the corresponding function. Firm-size quartiles are defined using annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. A request is counted once regardless of the number of downstream outputs or automated backend executions it triggers.

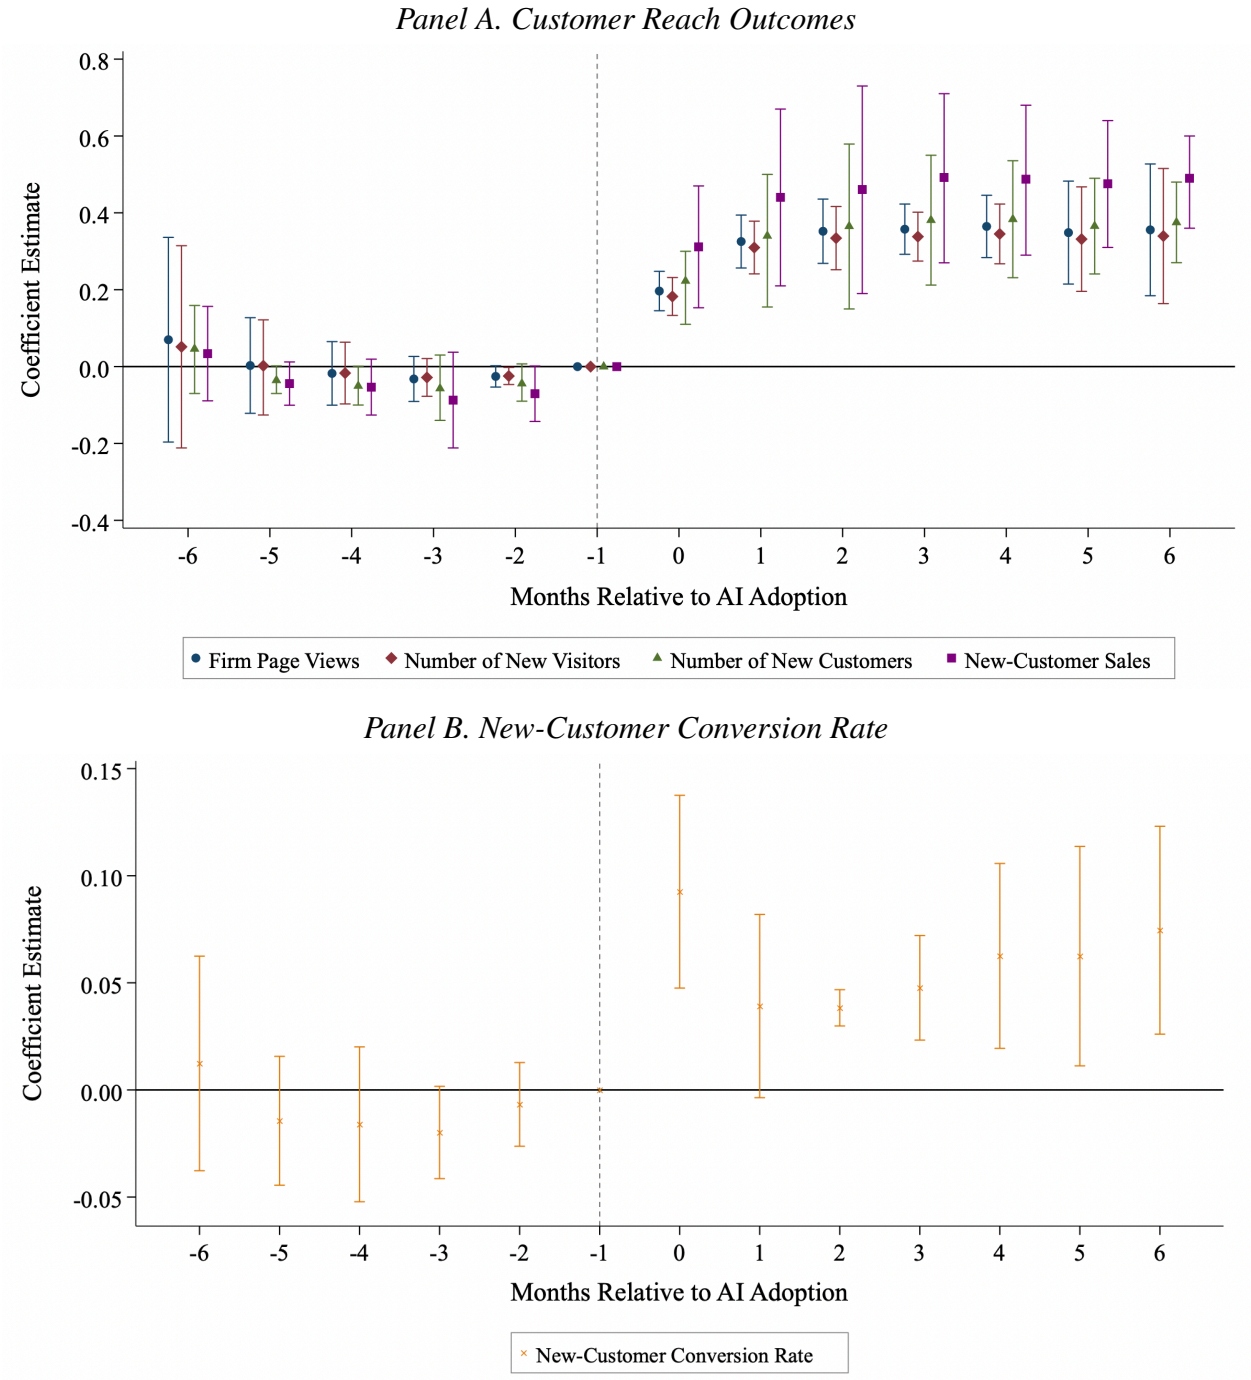
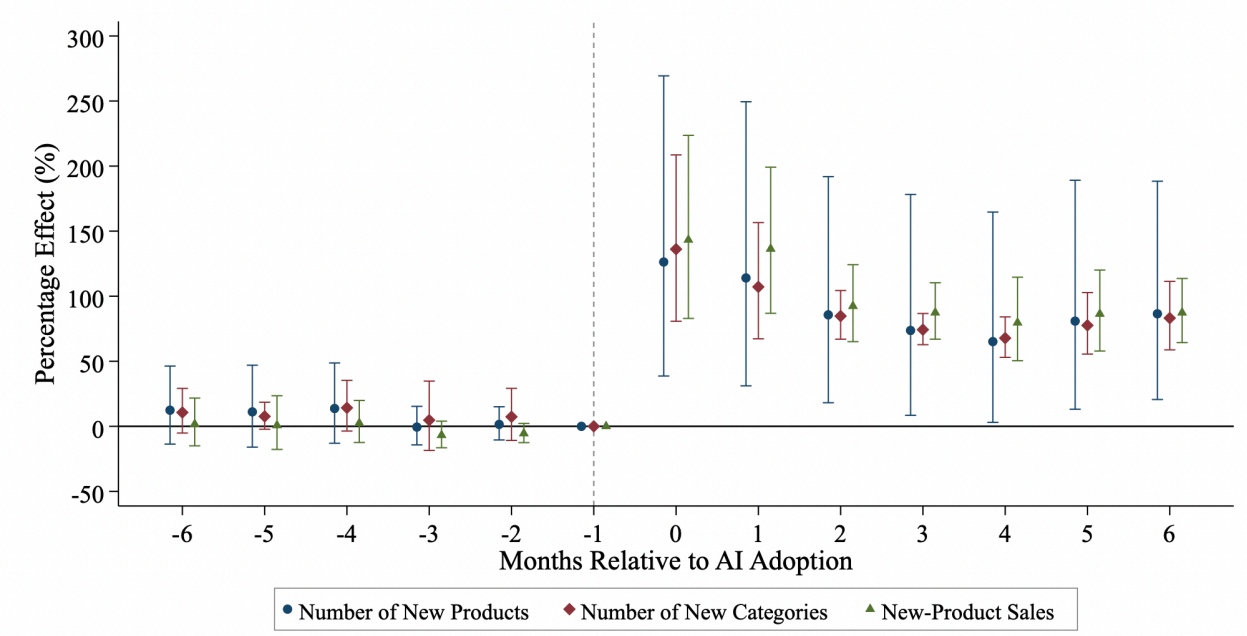


Figure D.1
Dynamic Effects of AI Adoption on Customer Reach

Notes: This figure reports balanced event-study estimates of the effects of AI adoption on customer reach. Panel A reports Firm Page Views, New Visitors, New Customers, and New-Customer Sales, all transformed using $\log(1+x)$. Panel B reports New-Customer Conversion Rate. New-Customer Conversion Rate is the share of first-time visitors who make their first purchase from the firm within the same month and is normalized by its matched-control pre-adoption mean. The event study uses the seven cohorts adopting from July 2024 through January 2025 and a common six-month window before and after adoption. Event month -1 is omitted, so each point is relative to the month immediately before adoption. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical bars denote 95% confidence intervals based on standard errors clustered at the Level-2 category level.

Panel A. Product Introductions and Sales



Panel B. New-Product Success Rate

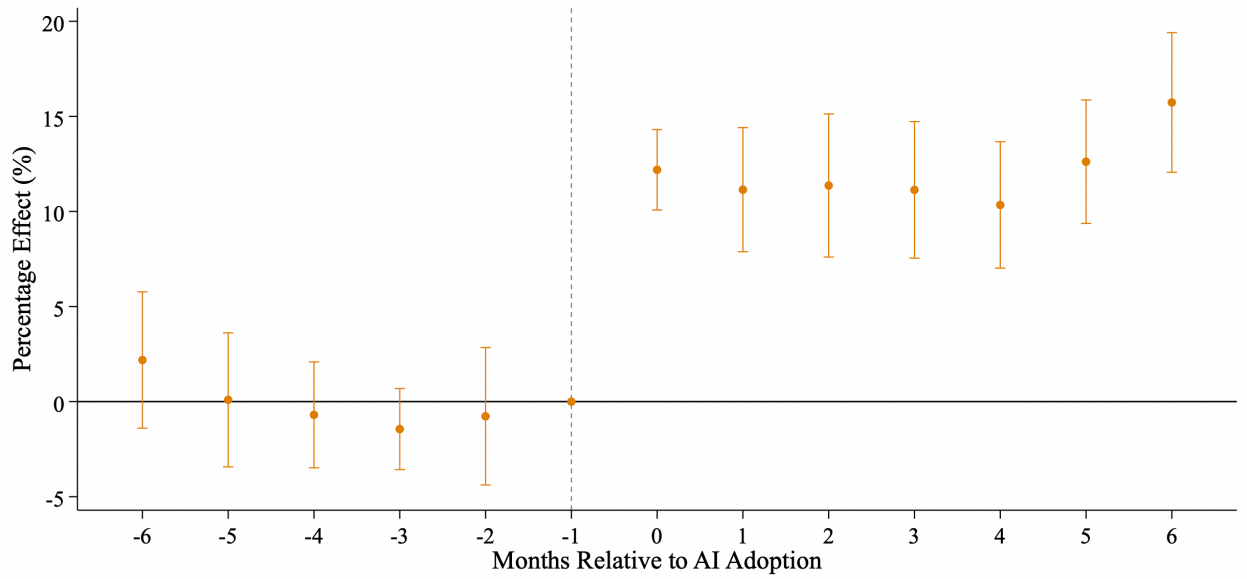
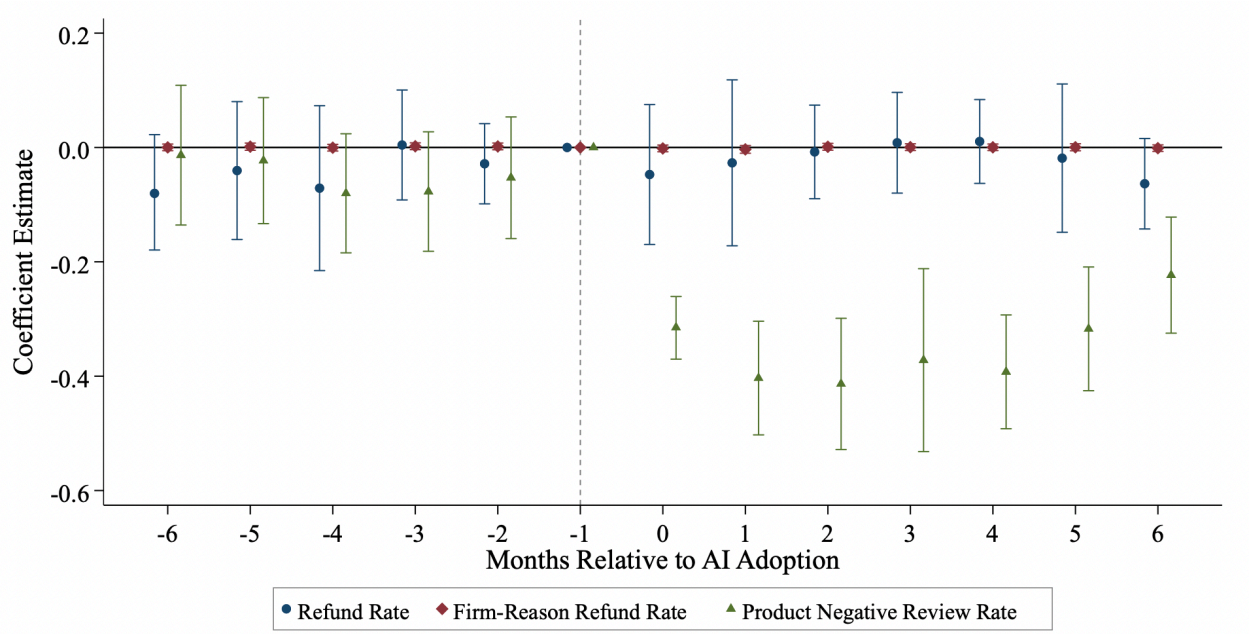


Figure D.2
Dynamic Effects of AI Adoption on Product Offerings

Notes: This figure reports balanced event-study estimates of the effects of AI adoption on product offerings. Panel A reports PPML estimates for the number of new products, number of new categories, and new-product sales, with percentage effects calculated as $100[\exp(\hat{\beta}) - 1]$. Panel B reports linear DID estimates for New-Product Success Rate, defined as the share of newly introduced products generating at least one sale. This outcome is observed only in firm-months with new introductions and normalized by the matched-control pre-adoption mean, with percentage effects calculated as $100\hat{\beta}$. The event study uses the seven cohorts adopting from July 2024 through January 2025 and a common six-month window before and after adoption. Event month -1 is omitted. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical bars denote 95% confidence intervals based on standard errors clustered at the Level-2 category level.

Panel A. Product Quality



Panel B. Service and Fulfillment Quality

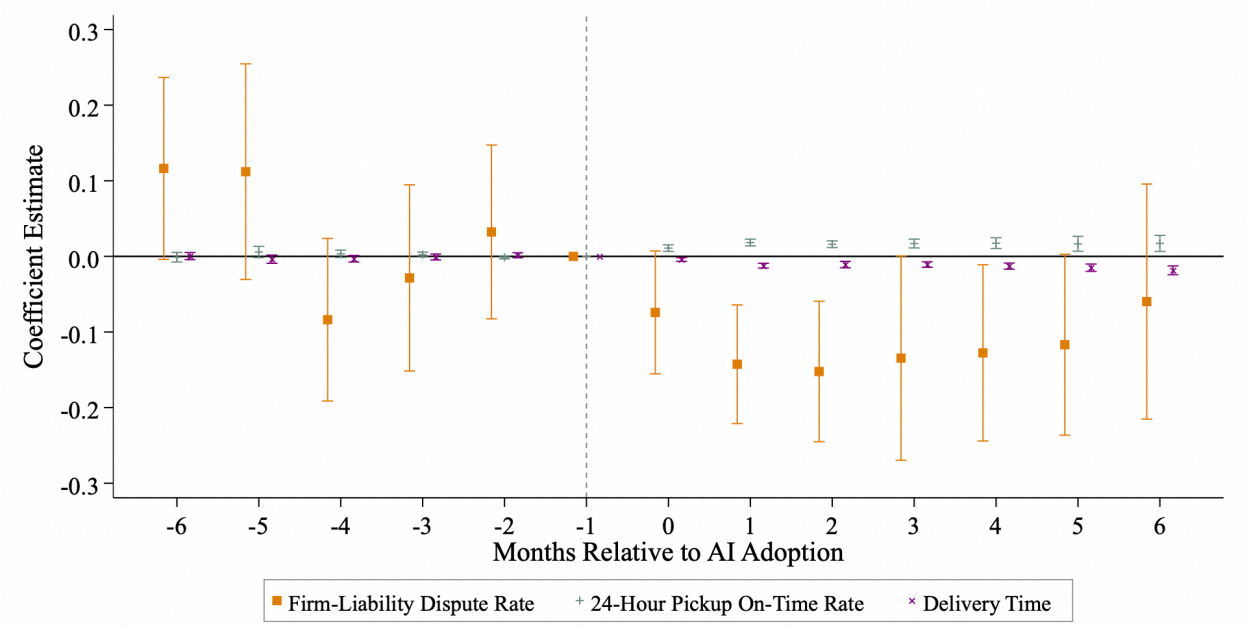


Figure D.3
Dynamic Effects of AI Adoption on Product and Service Quality

Notes: This figure reports balanced event-study estimates of the effects of AI adoption on product, service, and fulfillment quality. Panel A reports Refund Rate, Firm-Reason Refund Rate, and Product Negative Review Rate; Panel B reports Firm-Liability Dispute Rate, 24-Hour Pickup On-Time Rate, and Delivery Time. All rate outcomes are normalized by their matched-control pre-adoption means, while Delivery Time is transformed using $\log(1+x)$. The event study uses the seven cohorts adopting from July 2024 through January 2025 and a common six-month window before and after adoption. Event month -1 is omitted. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical bars denote 95% confidence intervals based on standard errors clustered at the Level-2 category level.

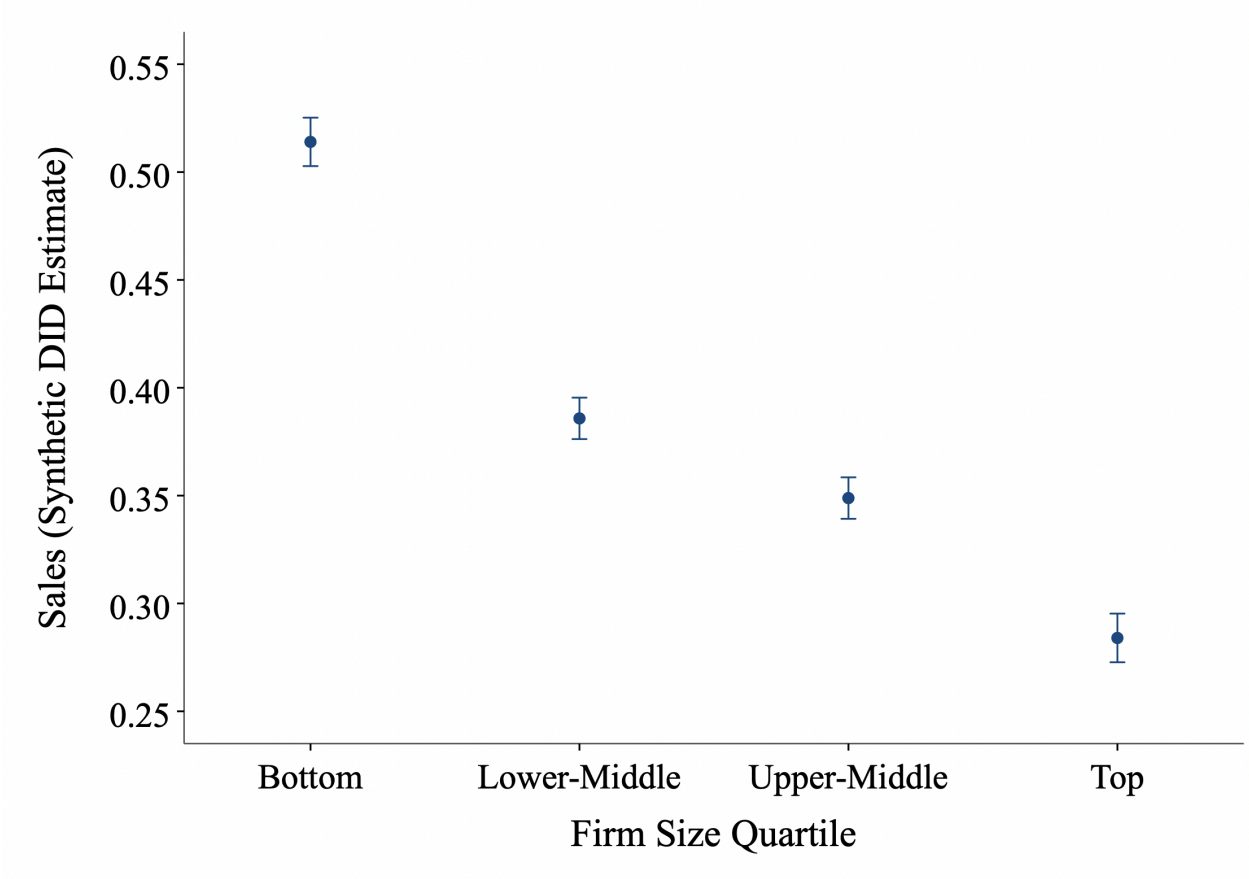


Figure D.4
Synthetic Difference-in-Differences Estimates by Firm Size

Notes: This figure reports synthetic difference-in-differences (SDID) estimates of the effect of AI adoption on monthly firm sales by pre-treatment firm-size quartile. The outcome is transformed using $\log(1 + x)$. SDID is estimated separately within cells defined by Level-1 product category, firm-size quartile, and adoption cohort (July 2024–June 2025). Within each size quartile, cell-specific estimates are averaged using the number of treated firms as weights. Firm-size quartiles are based on pre-treatment annual sales, with cutoffs at RMB 24,460, RMB 63,088, and RMB 147,716. Dots indicate estimated effects in log points, and vertical lines denote 95% confidence intervals.

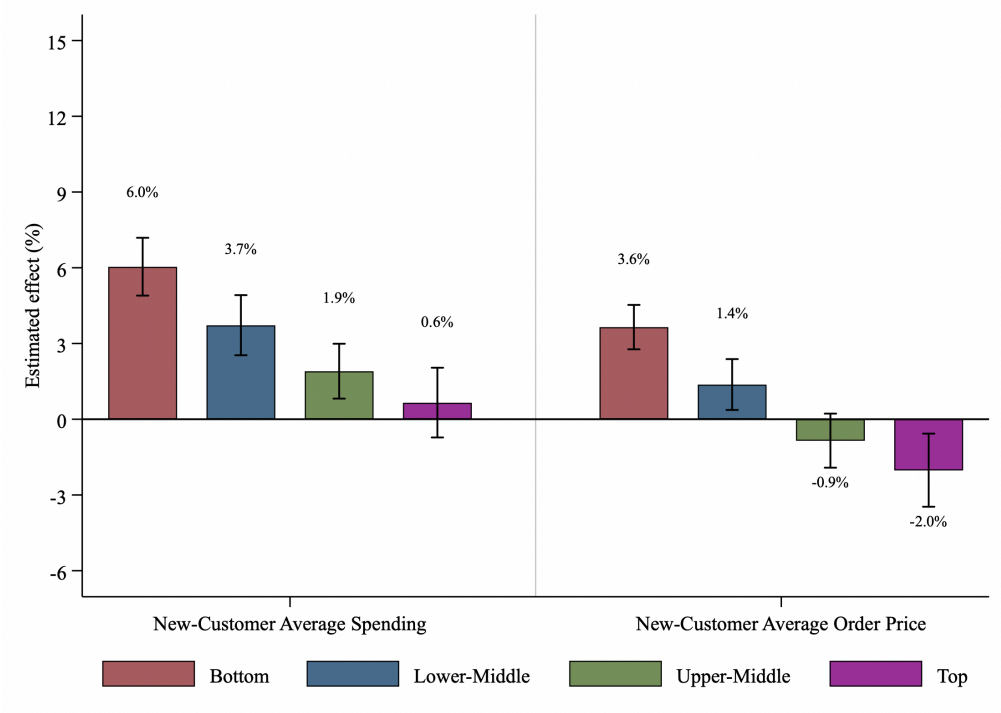


Figure E.1
Heterogeneous Effects of AI Adoption on New-Customer Value

Notes: This figure reports matched difference-in-differences estimates of AI adoption on the value of newly acquired customers by pre-treatment firm-size quartile. The outcomes are average spending per new customer and average order price for new-customer transactions, both transformed using $\log(1+x)$. New customers are consumers making their first-ever purchase from the firm, identified using the platform's full historical records. Bars report percentage effects calculated as $100[\exp(\hat{\beta})-1]$, and vertical lines denote 95% confidence intervals. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects, with standard errors clustered at the Level-2 category level.

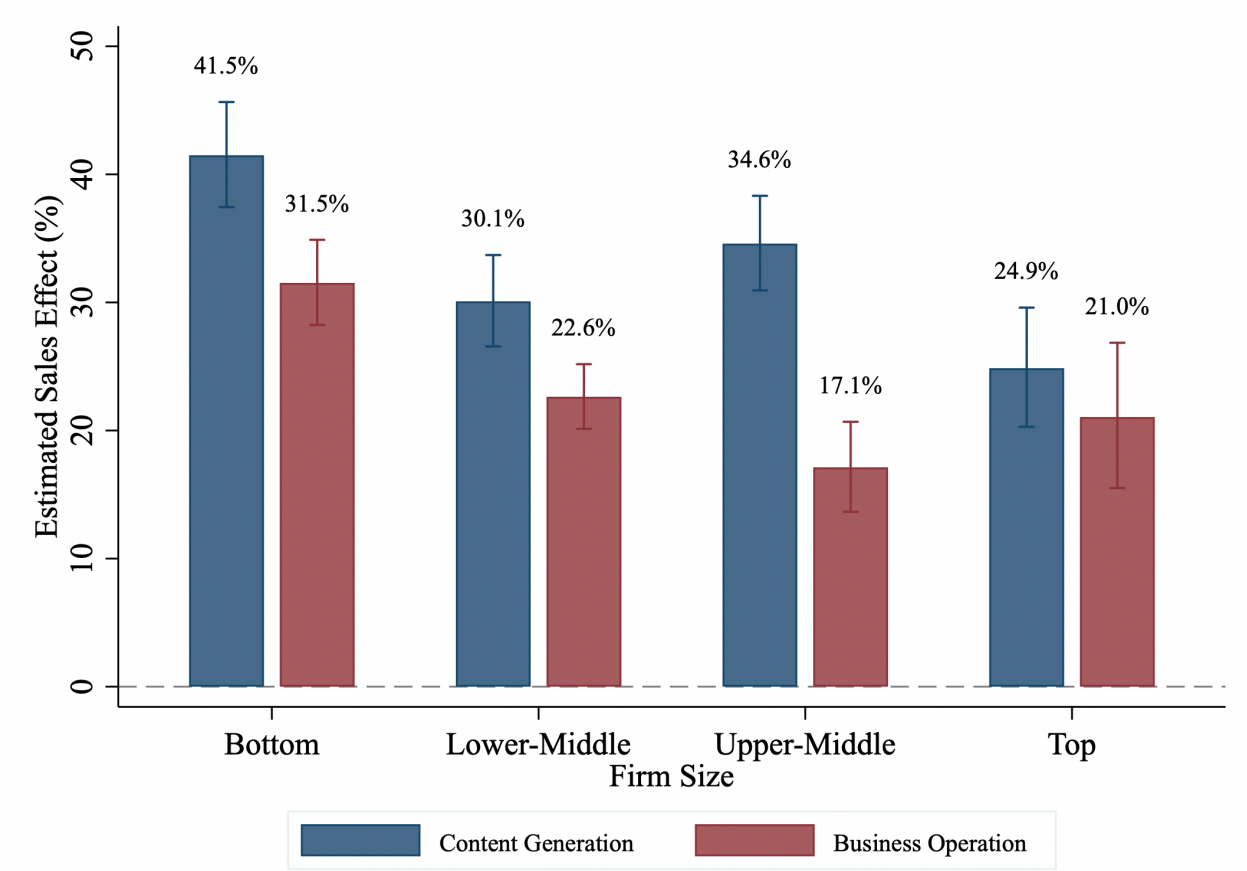


Figure E.2
Heterogeneous Effects by Firm Size: Single AI Functionalities

Notes: This figure reports matched difference-in-differences estimates of the effect of adopting each AI functionality separately on monthly sales. Content Generation Only firms adopt Content Generation but never Business Operation; firms first adopting Content Generation before July 2024 are excluded to align the cohort window with the main analysis. Business Operation Only firms adopt Business Operation but never Content Generation. Control firms adopt neither functionality. Bars report percentage effects on monthly sales, calculated as $100[\exp(\hat{\beta}) - 1]$, and vertical lines denote 95% confidence intervals. Firm-size groups use the baseline annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects, with standard errors clustered at the Level-2 category level.

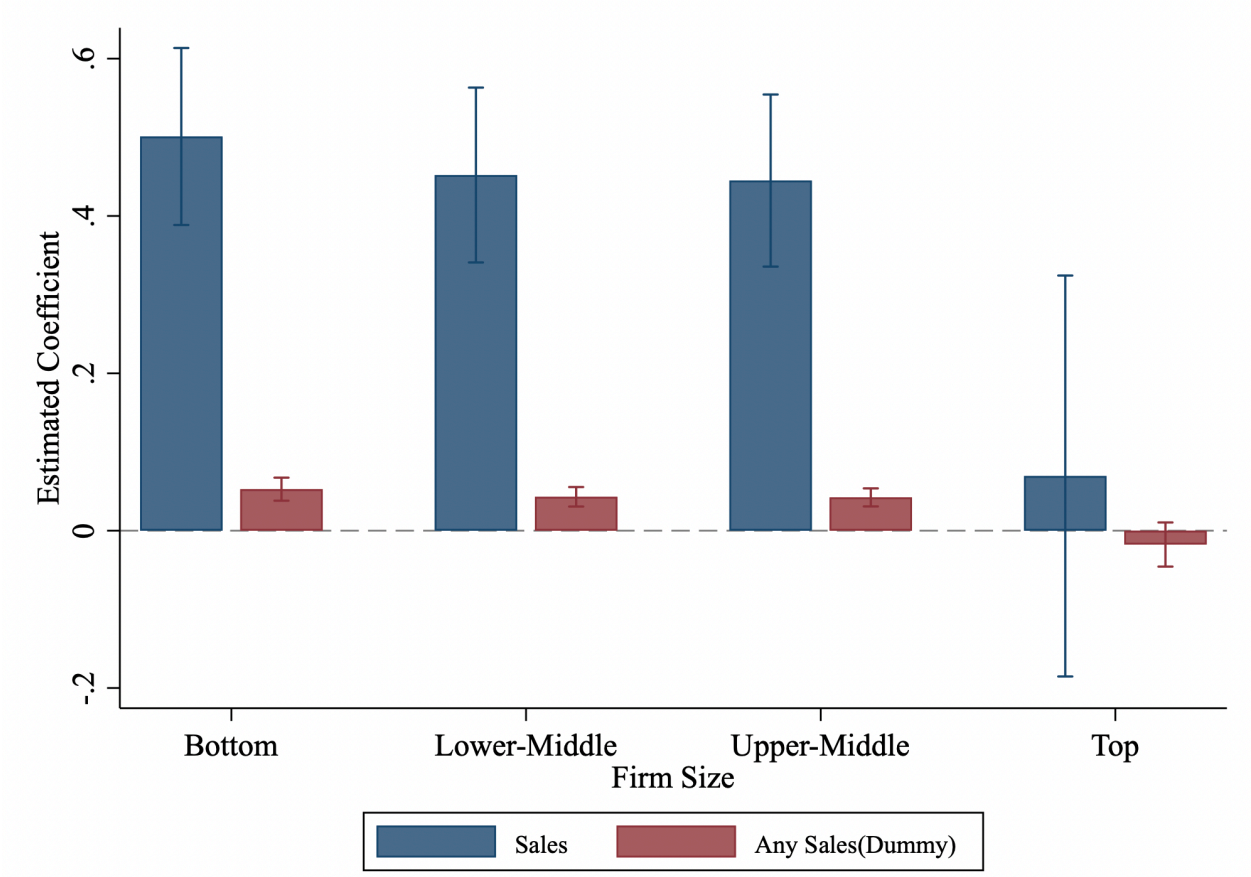


Figure E.3
Heterogeneous Effects by Firm Size: Allowing Zero-Sales Months

Notes: This figure reports matched difference-in-differences estimates by firm size in the expanded sample that permits zero-sales months; all other treatment, control, and matching definitions follow the baseline analysis. Sales is transformed using $\log(1+x)$. Any Sales equals one if the firm records positive sales in a given month and zero otherwise. Firm-size groups are quartiles of pre-treatment annual sales in the expanded sample, with cutoffs of RMB 13,126, RMB 30,259, and RMB 58,632. Bars report estimated treatment effects and vertical lines denote 95% confidence intervals. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects, with standard errors clustered at the Level-2 category level.

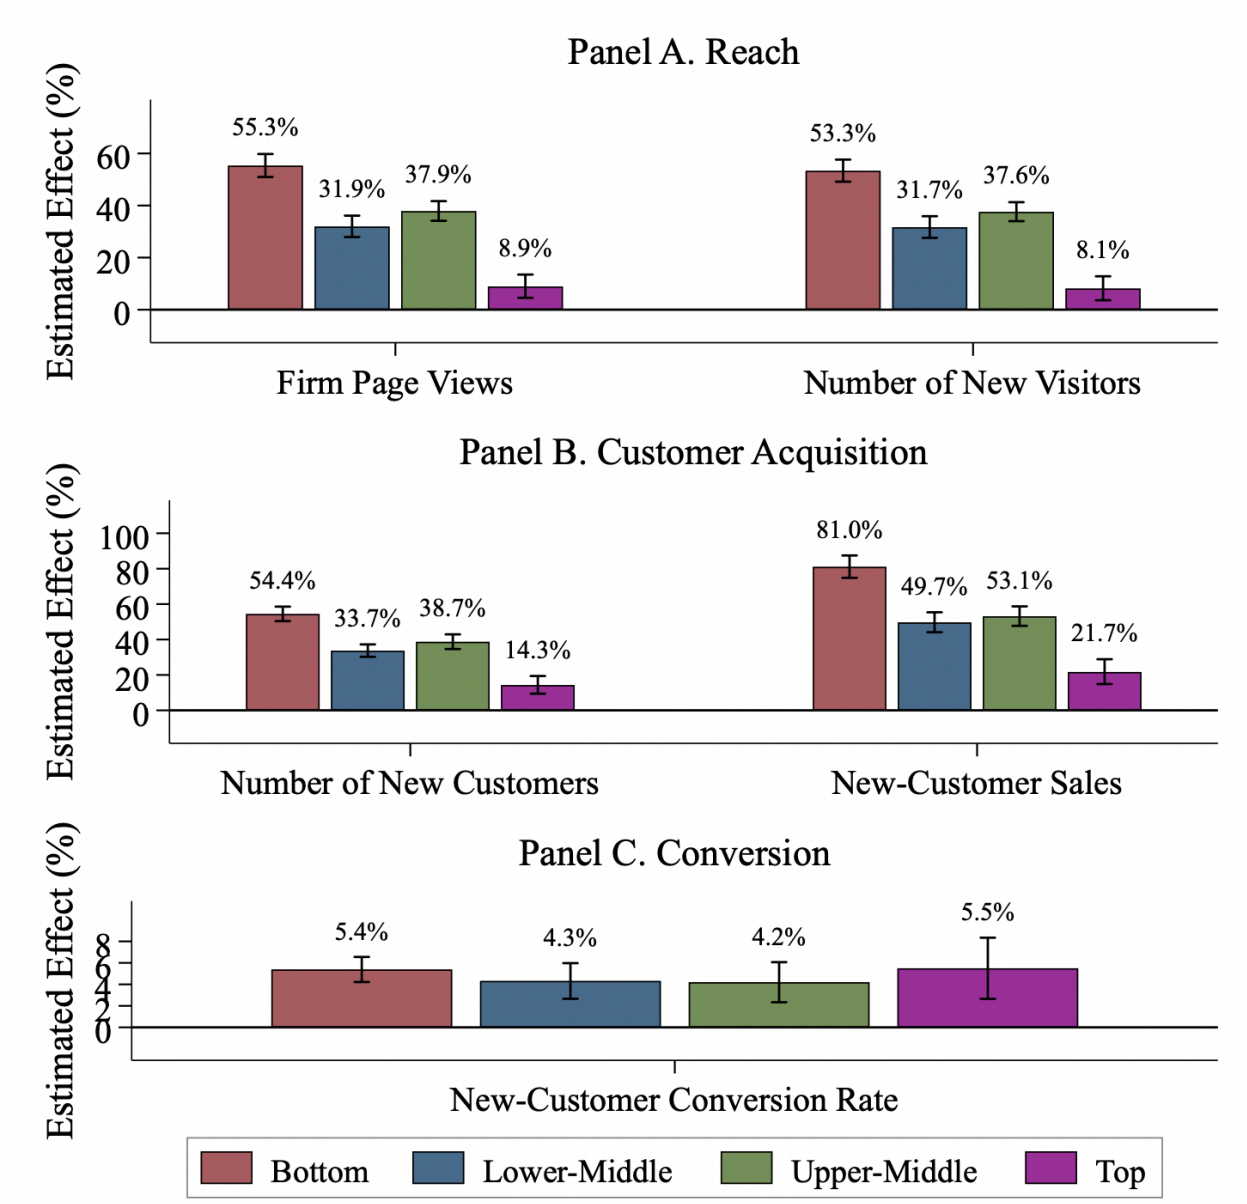


Figure E.4
Heterogeneous Effects of AI Adoption on Customer Reach

Notes: This figure reports matched difference-in-differences estimates of AI adoption on customer reach by pre-treatment firm-size quartile. Panel A reports Firm Page Views and New Visitors; Panel B reports New Customers and New-Customer Sales; Panel C reports New-Customer Conversion Rate. The first four outcomes are transformed using $\log(1 + x)$. New-Customer Conversion Rate is the share of first-time visitors who make their first purchase from the firm within the same month and is normalized by its matched-control pre-adoption mean. Bars report percentage effects: $100[\exp(\hat{\beta}) - 1]$ for log-transformed outcomes and $100\hat{\beta}$ for the normalized conversion rate. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical lines denote 95% confidence intervals based on standard errors clustered at the Level-2 product-category level.

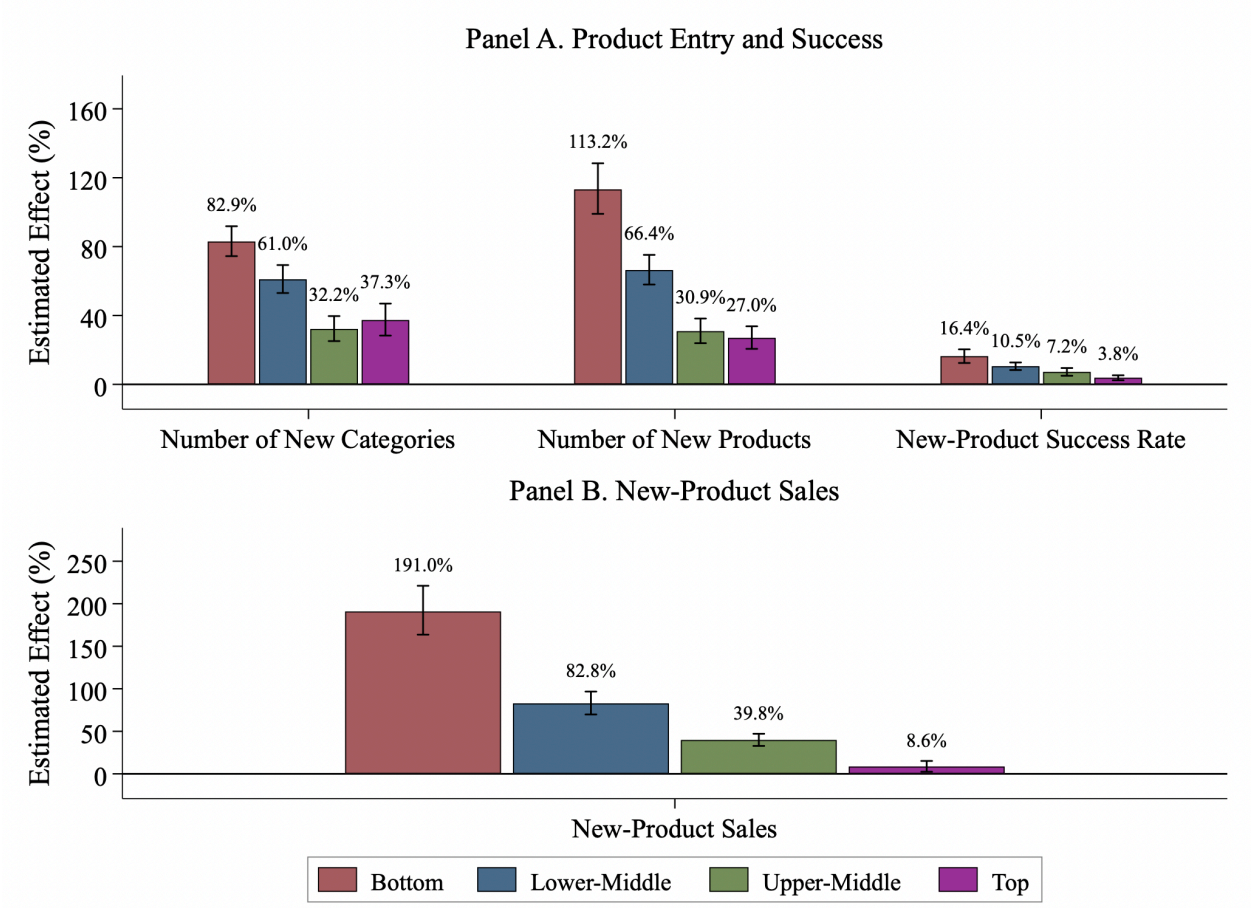


Figure E.5
Heterogeneous Effects of AI Adoption on Product Offerings

Notes: This figure reports matched DID estimates of AI adoption on product-line adaptation by pre-adoption firm-size quartile. Panel A reports new categories, new products, and new-product success rates; Panel B reports new-product sales. All outcomes except New-Product Success Rate are estimated using PPML, with outcomes winsorized at the 99th percentile. New-Product Success Rate is the share of new products generating at least one sale, observed only in firm-months with new introductions and estimated using linear DID after normalization by its matched-control pre-adoption mean. Bars report percentage effects calculated as $100[\exp(\hat{\beta}) - 1]$ for PPML estimates and $100\hat{\beta}$ for the normalized success rate. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, 63,088, and 147,716. All regressions include firm-by-cohort and cohort-by-month fixed effects and matched-sample weights. Vertical lines denote 95% confidence intervals based on standard errors clustered at the Level-2 product-category level.

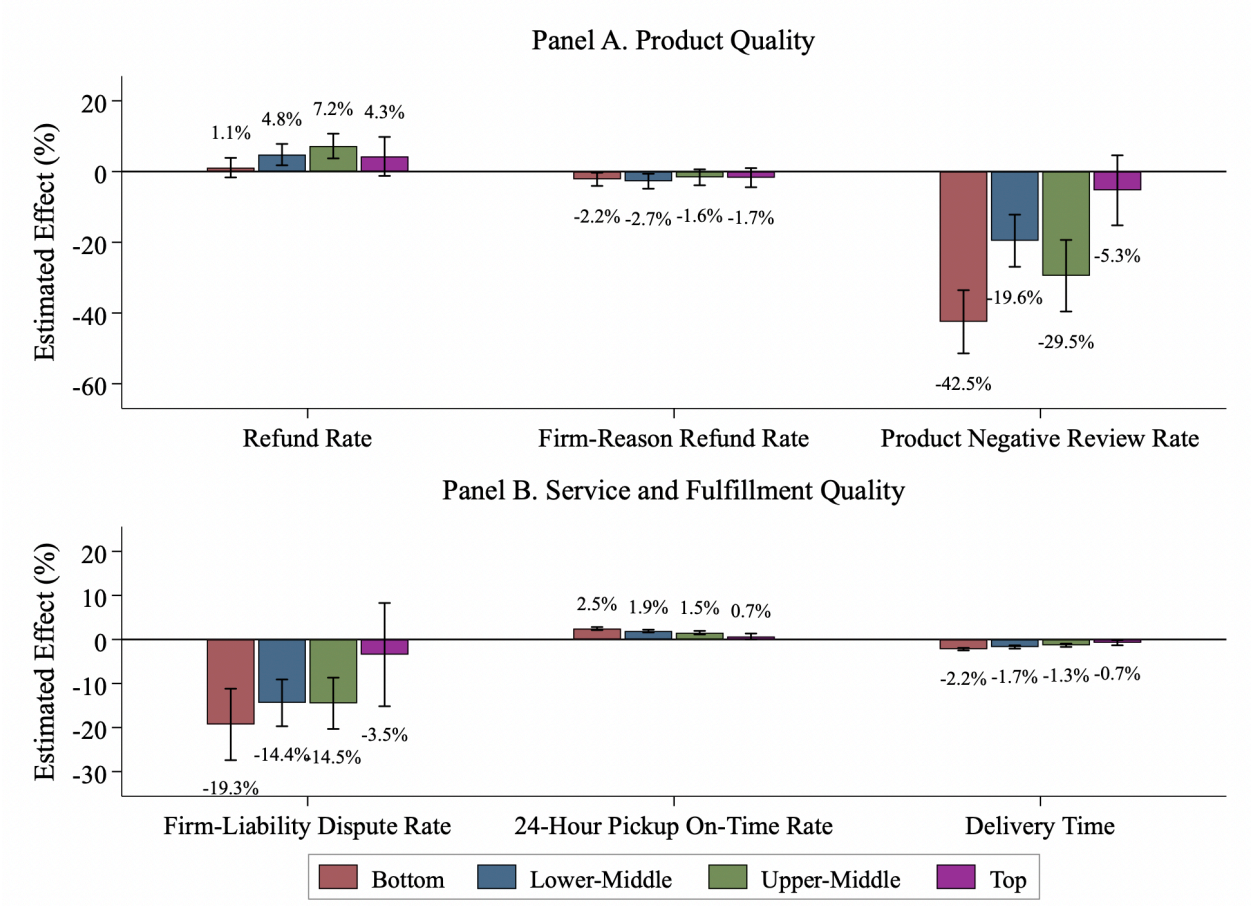


Figure E.6
Heterogeneous Effects of AI Adoption on Product and Service Quality

Notes: This figure reports matched difference-in-differences estimates of AI adoption on product, service, and fulfillment quality by pre-treatment firm-size quartile. Panel A reports Refund Rate, Firm-Reason Refund Rate, and Product Negative Review Rate; Panel B reports Firm-Liability Dispute Rate, 24-Hour Pickup On-Time Rate, and Delivery Time. The rate outcomes are normalized by their matched-control pre-adoption means; Delivery Time is transformed using $\log(1+x)$. Bars report percentage effects as $100\hat{\beta}$ for normalized rate outcomes and $100[\exp(\hat{\beta})-1]$ for Delivery Time. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical lines denote 95% confidence intervals based on standard errors clustered at the Level-2 product-category level.

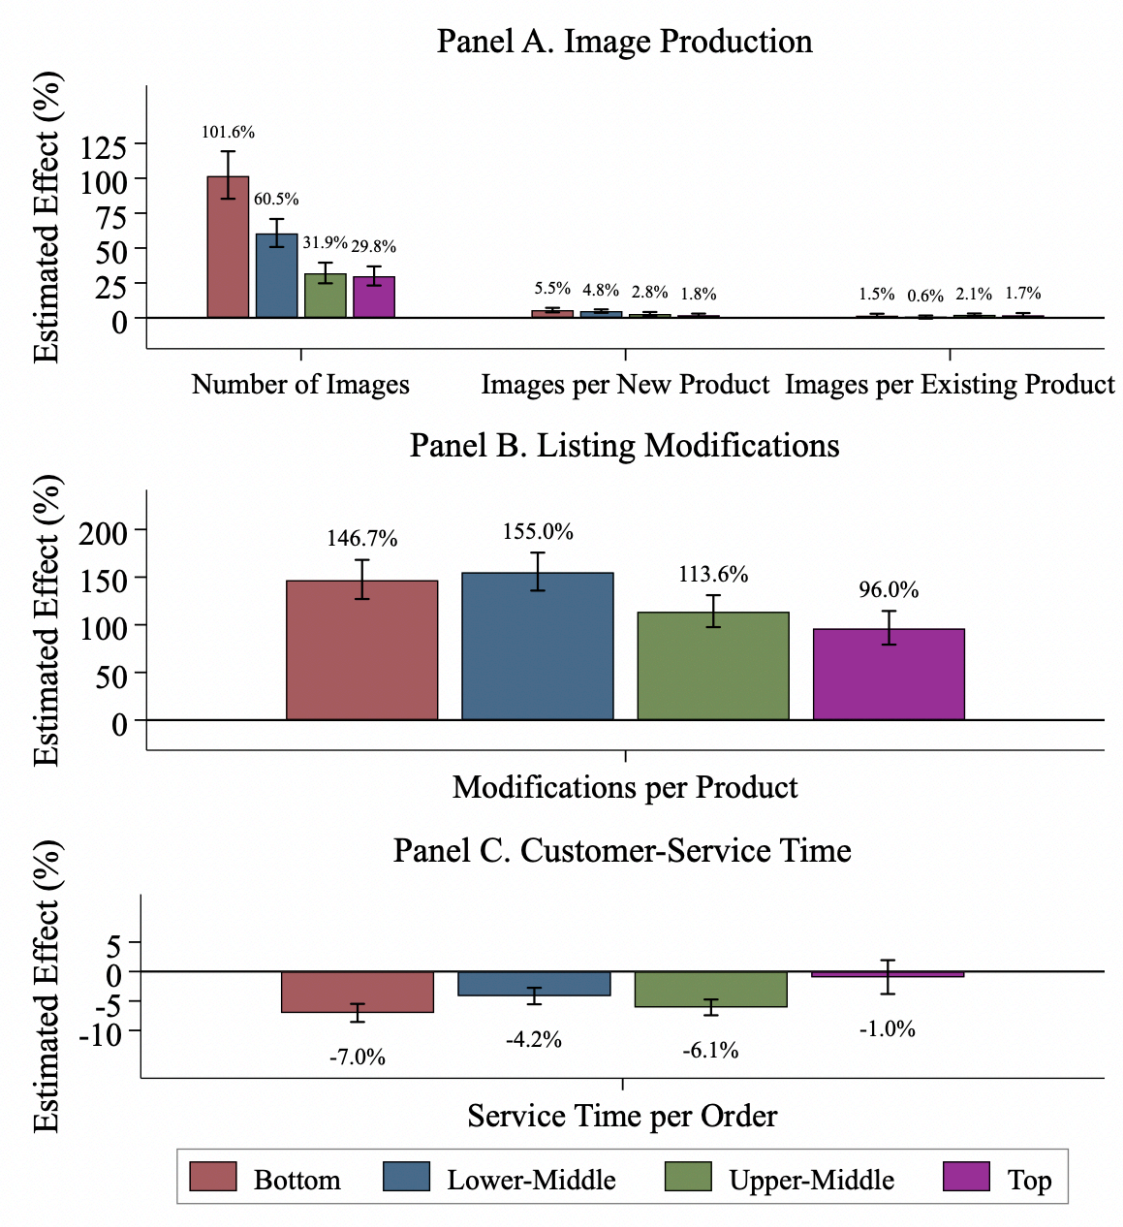


Figure E.7
Heterogeneous Effects of AI Adoption on Operating Activities

Notes: This figure reports matched DID estimates of AI adoption on operating activities and inputs by pre-adoption firm-size quartile. Panels A–C report image production, listing modifications, and customer-service activities, respectively. Number of Images and Modifications per Product are estimated using PPML with outcomes winsorized at the 99th percentile; the remaining outcomes use linear DID with $\log(1+x)$ transformations. Bars report $100[\exp(\hat{\beta})-1]$, measuring proportional changes in x for PPML estimates and in $1+x$ otherwise. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, 63,088, and 147,716. All regressions include firm-by-cohort and cohort-by-month fixed effects and matched-sample weights. Vertical lines denote transformed 95% confidence intervals based on standard errors clustered at the Level-2 product-category level.

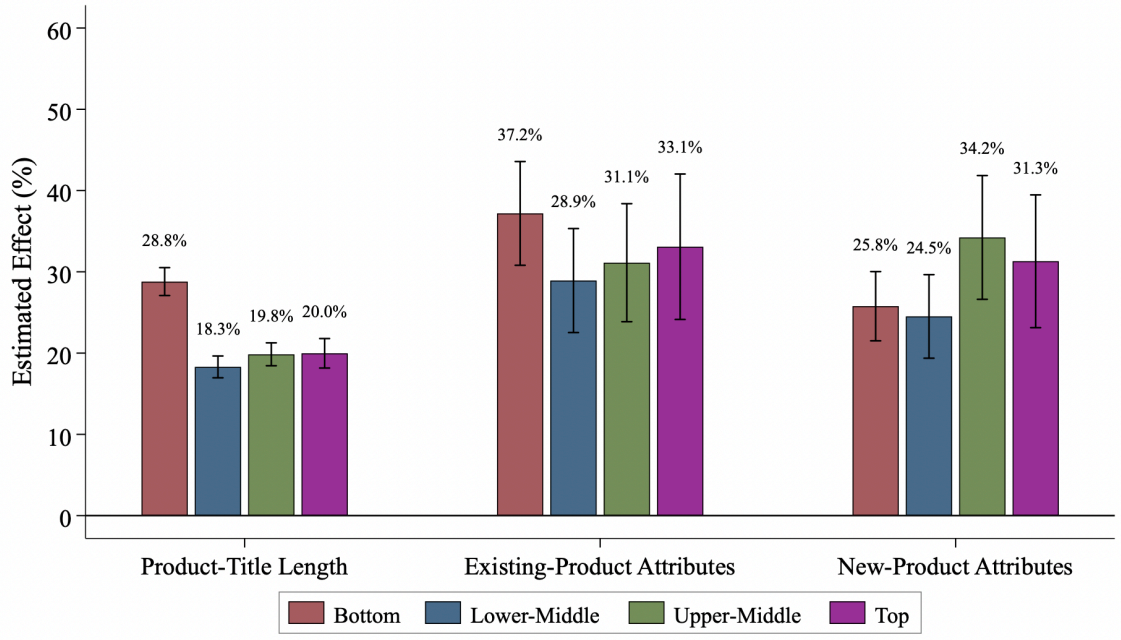


Figure E.8
Heterogeneous Effects of AI Adoption on Product Information Provision

Notes: This figure reports matched difference-in-differences estimates of AI adoption on product information provision by pre-launch firm-size quartile. Product-Title Length is transformed using $\log(1+x)$; its bars report percentage effects calculated as $100[\exp(\hat{\beta})-1]$. Existing- and New-Product Attribute Coverage measure filled product attributes relative to the category-specific maximum and are estimated in levels; their bars report the coefficient relative to the corresponding matched-control pre-adoption mean, $100\hat{\beta}/\bar{Y}_{0,C}$. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Vertical lines denote 95% confidence intervals based on standard errors clustered at the Level-2 product-category level.

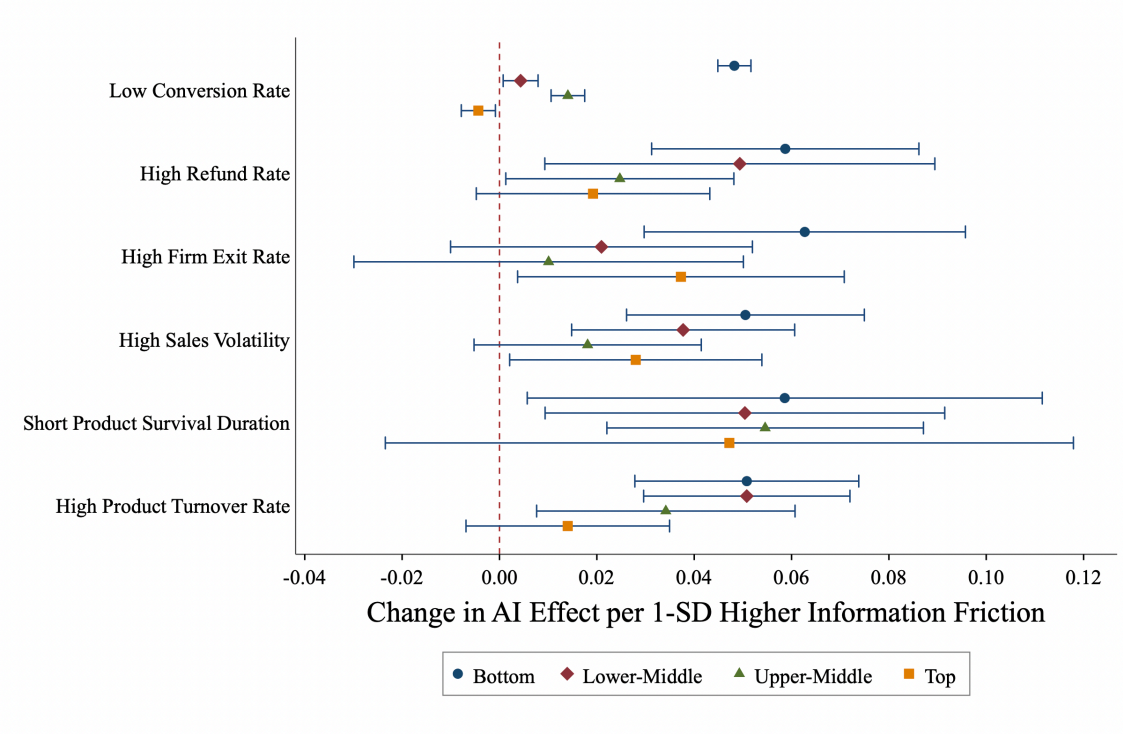


Figure E.9
Heterogeneous Sales Effects by Baseline Information Frictions

Notes: This figure reports how baseline information frictions modify the effect of AI adoption on monthly sales within each pre-launch firm-size quartile. The dependent variable is log(1 + Sales), and the six predetermined information-friction measures are those defined in Table A.3. Conversion Rate and Product Survival Duration are multiplied by −1 before standardization so that, for every measure shown, a higher value denotes greater information frictions. Each point is the coefficient on the interaction between Treat × Post and the standardized friction measure; it therefore gives the change in the AI treatment effect associated with a one-standard-deviation increase in information frictions. All specifications include the relevant lower-order interaction terms, with coefficients omitted for brevity. Firm-size quartiles use annual-sales cutoffs of RMB 24,460, RMB 63,088, and RMB 147,716. Regressions include firm-by-cohort and cohort-by-calendar-month fixed effects. Horizontal lines denote 95% confidence intervals based on standard errors clustered at the Level-2 category level.