Social Media, Information Networks, and International Trade*

George Cui
International Monetary Fund

Kailin Gao
University of International Business and Economics

March 2026

Preliminary and incomplete. Please do not circulate.

Abstract

We study how a professional digital platform reduces information frictions in international trade. Combining platform data on more than 20 million firm-to-firm online interactions with Chinese customs records and input–output relationships, we establish four facts: (i) platform adoption raises firm exports by 12 percent within two years, operating entirely through the extensive margin; (ii) adoption is unequal across regions and firms, with richer cities and larger firms adopting disproportionately; (iii) information flows remain geographically local even in a digital environment; (iv) knowledge transmission concentrates along production networks, flowing from upstream suppliers to downstream buyers. We develop a multiregion, multisector trade model with variable-markup oligopolistic competition in which location-sector owners endogenously post knowledge that builds information stocks through domestic production networks.

JEL codes: F12, F14, L15, O33
Keywords: Information frictions, digital platforms, international trade, variable markups, knowledge diffusion

*We are grateful to Kathryn Dominguez, James Hines, Andrei Levchenko, Sebastian Sotelo, and Petia Topalova for their invaluable guidance and support. We also thank Yan Bai, Yan Chen, Jaedo Choi, Nan Li, Tanya Rosenblat, Younghun Shim and seminar participants at the University of Michigan Economics Department and Information School, Peking University China Center of Economic Research (CCER) Summer Institute, and the Chinese Economists Society (CES) North America Annual Conference 2025 for their insightful comments. Cui gratefully acknowledges support from the Macroeconomic Research on Climate Change and Emerging Risks in Asia program of the Ministry of Economy and Finance of the Government of Korea, as well as financial support from the Ross China Initiative. Gao gratefully acknowledges funding from the National Natural Science Foundation of China (Grant No. 72103042), and support from the Scientific Research Laboratory of Digital Trade at the University of International Business and Economics. All errors are ours. The views in this paper are those of the authors and do not necessarily reflect the views of the International Monetary Fund, its Executive Board, and IMF Management. This paper was previously circulated under the title "Information, Social Media and International Trade: Theory and Evidence Using Twenty Million Online Postings." Email: gcui@imf.org, and kailingao@uibe.edu.cn.

1. Introduction

Every day, billions of individuals post, share, and consume digital content on social media platforms. Social media has revolutionized information dissemination, enabling individuals to learn from others across vast distances in real time. The economic significance of social media is substantial: in 2023, the top 10 social media companies had a combined market value of $6 trillion, representing 12 percent of total U.S. market capitalization.[1] Beyond their financial footprint, social media platforms influence economic behavior by shaping consumer preferences, investment decisions, and market trends.[2] However, estimating the economic impact of social media poses several empirical challenges: micro-level measures of adoption and treatment effects are rare, users' economic activities are difficult to quantify, and misinformation can distort decision-making.[3]

Information frictions are central to international trade. Exporting to foreign markets is widely recognized as an information-intensive activity, requiring firms to gather and process substantial knowledge about consumer preferences, competitive landscapes, distribution networks, regulatory environments, and market-specific tacit know-how (Rhee, 1989; Rauch, 2001; Rauch and Watson, 2002; Theodosiou and Katsikea, 2013).[4] The existence of information frictions and associated externalities justifies export promotion policies (Hausmann and Rodrik, 2003; Volpe Martincus and Carballo, 2012; Wei, Wei and Xu, 2021). Yet there is little empirical evidence of how firms actually acquire the information they need to export. At the same time, when trade policy shifts rapidly or geopolitical fragmentation makes established business routines less reliable, these frictions become especially salient.

This paper fills the empirical gap by leveraging novel data on international trade. We assemble platform records from a large Chinese professional social media site devoted to international trade and link them to firm-level customs transactions and domestic production-

[1] Source: GlobalData. Goldman Sachs estimates the creator economy will grow to $480 billion by 2027 (Goldman Sachs, 2023).

[2] Bailey et al. (2022) study peer effects on Facebook. Antweiler and Frank (2004) find that online stock messages help predict market volatility. Bollen, Mao and Zeng (2011) show that collective mood states measured on Twitter correlate with the Dow Jones Industrial Average.

[3] Allcott and Gentzkow (2017) examine misinformation on social media during the 2016 U.S. election.

[4] Rhee (1989) study how Hong Kong trading companies deepen export marketing and find that the most valuable resource exporters have is deep knowledge about external markets, buyers, and local production capabilities.

network data. The platform's public discussion format allows us to observe who adopts, who posts, where information flows, and how those flows align with sectoral and regional economic structures. For example, a typical thread involves one user asking how to navigate regulatory changes affecting exports to a specific destination, and another user responding with detailed guidance on compliance strategies, alternative routing, and buyer-search methods. Merging the online dataset with the universe of transaction-level Chinese Customs Data by firm names, we find that at least 10 percent of Chinese exporting firms from 2000 to 2017 had adopted the platform.

Using this comprehensive dataset, we document four new empirical facts.

Fact 1: Platform adoption raises firm-level exports. Using a stacked event-study design with propensity-score matching, we find that adopting firms increase their exports by roughly 12 percent within two years. The gains come entirely from selling to new destination markets—not from higher prices or larger shipments to existing buyers. Moreover, the returns to adoption are larger where platform information is already denser, consistent with external economies of scale in knowledge.

Fact 2: Adoption is uneven across regions and firms. Despite being freely accessible online, the platform is not adopted uniformly. Adoption rates are systematically higher in richer cities with better internet infrastructure, and larger firms are far more likely to adopt—adoption rates rise 12-fold from the bottom to the top export-size quintile. These patterns suggest that effective use of the platform requires complementary capabilities that are unevenly distributed.

Fact 3: Information flows remain geographically local. Even in a digital environment where communication costs are geographically invariant, information flows on the platform decline sharply with physical distance. A 10 percent increase in distance reduces bilateral information flows by 6–9 percent. This finding echoes the classic result that knowledge spillovers are localized, and extends it to a setting where the technology of communication is explicitly distance-neutral.

Fact 4: Knowledge transmission concentrates along production networks. Information sharing is directional: it flows predominantly from upstream suppliers to downstream buyers along domestic production chains. This asymmetry supports an intentional-teaching mechanism in which suppliers share export knowledge with their buyers because the buyer's success

abroad generates future demand for the supplier's inputs.

On the theory side, we incorporate these information flows into an oligopolistic trade model in the spirit of Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015). In that environment, information matters not only because it lowers export-side marginal costs—it also affects residual demand elasticities, market shares, and markups. Firms draw heterogeneous productivities, choose destination entry, may pay a fixed cost to learn from a local information stock, and compete in destination markets with share-dependent elasticities. The information stock is endogenous: it is generated by location-sector owners who internalize local firm profits and therefore choose outbound knowledge flows strategically. This unified-owner formulation collapses the firm problem and the innovator problem into a single maximization, yielding a knowledge-posting condition that decomposes naturally into demand, market-share, and markup channels.

Introducing variable markups changes both the economics of the model and its quantitative implications. Under constant markups, information works mainly through costs and exporter cutoffs. Under oligopoly with variable markups, the same information shock alters competitive pressure inside each destination-sector market. Analytically, the knowledge-posting first-order condition becomes implicit rather than closed-form. Quantitatively, it implies that the resulting information buffer is more tightly disciplined by moments capturing concentration and competition.

The quantitative results show that the information mechanism is large under oligopolistic competition. In the main specification, a 40 percent tariff increase between China and the rest of the world reduces China's welfare by 32.35 percent when information flows are shut down, but by only 23.41 percent when they remain active, implying an information buffer of 8.94 percentage points. The same comparison yields an export buffer of 12.35 percentage points. Information thus offsets roughly one-quarter of the welfare cost of fragmentation, underscoring that a theory of trade and information should take both network structure and market structure seriously.

Related literature. Our paper contributes to three strands of work.

First, the economics of knowledge flows. There is a long-standing tradition of studying the economic importance of knowledge (Hayek, 1945). Measuring knowledge flows is inherently

challenging because of their intangible nature.[5] Jaffe, Trajtenberg and Henderson (1993) introduced patent citations as proxies for knowledge spillovers, and subsequent work has used citation data extensively (Caballero and Jaffe, 1993; Ellison, Glaeser and Kerr, 2010; Akcigit and Kerr, 2018; Fadeev, 2024; Ahn et al., 2026; Choi, 2024). A rising literature uses non-conventional data sources to measure information, including newspapers (Bui et al., 2022) and online platforms (Cotterlaz and Guillouzouic, 2025). We advance this literature by pioneering the use of online postings on a social media platform to measure firm-to-firm knowledge flows directly. Unlike patent citations, online postings measure knowledge that is rarely patented but economically important—information about foreign market fundamentals and practical export know-how. Our data provide a direct measure of learning through social interactions, which connects to the theoretical literature on idea diffusion (Alvarez et al., 2013; Perla and Tonetti, 2014; Lucas Jr and Moll, 2014; Sampson, 2016; Buera and Lucas Jr, 2018; Buera and Oberfield, 2020; Perla, Tonetti and Waugh, 2021).

An important feature of our data is that online postings are direct evidence of intentional knowledge transfer, shedding light on the mechanism behind knowledge flows. Understanding this mechanism is vital for endogenous growth, international trade (Jones, 2005; Aghion, Akcigit and Howitt, 2014; Melitz and Redding, 2021), and economic agglomeration (Marshall, 2009; Carlino and Kerr, 2015). Traditionally, knowledge spillovers are assumed to occur unintentionally, creating a free-rider problem and justifying government intervention (Bloom, Van Reenen and Williams, 2019; Bartelme et al., 2025). Fadeev (2024) find that patent citations concentrate among business partners, suggesting intentional cooperation rather than passive spillover. Our paper shows that intentional knowledge sharing exists between firms engaged in international trade and provides a mechanism: firms rationally help their downstream buyers to export so as to improve their own profits through supply-chain linkages.

Second, information frictions in international trade. A central question is why international trade remains significantly lower than domestic trade even after accounting for geographic distance. One explanation is substantial bilateral trade frictions related to culture, institutions, and information (Rauch, 2001; Anderson and Van Wincoop, 2004; Allen, 2014).

[5] Krugman (1991) notes that knowledge flows "are invisible; they leave no paper trail by which they may be measured and tracked."

Reducing these frictions could lead to significant trade increases and welfare improvements. Empirical evidence on information frictions often stems from studies on incomplete information about prices of homogeneous goods. Direct measurement of the information diffusion process is rare. A small literature connects observable communication to aggregate trade flows using travel (Cristea, 2011; Bernard, Moxnes and Saito, 2019; Startz, 2021; Tian and Yu, 2023), telephone calls (Portes and Rey, 2005), and website viewer history (Carballo et al., 2022). Using online postings, we observe communication and outcomes at the micro level, allowing us to build a fully specified model of firm behavior with underlying heterogeneity of information frictions. Our theory builds on the production-networks literature (Carvalho, 2014; Carvalho and Tahbaz-Salehi, 2019; Liu, 2019; Acemoglu and Azar, 2020; Arkolakis, Huneeus and Miyauchi, 2023), showing that the architecture of production networks facilitates collaborative behavior among firms. Hajdini et al. (2025) show that supply chain linkages are a key channel through which macroeconomic information propagates across firms, complementing our finding that network position shapes incentives to share export know-how.

Third, quantitative trade models with endogenous markups and technology adoption. We embed a network-based information mechanism inside the finite-firm environment of Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015). This complements recent work on technology adoption in trade (Choi and Shim, 2022; Buera and Oberfield, 2020; Bartelme et al., 2025) and the growing literature on the economic impact of information and communication technology (ICT) (Rosenblat and Mobius, 2004). While it is widely believed that ICT affects economic growth and inequality, empirical evidence is scant. The rich customs data enable us to evaluate the real economic effects of digital public goods, and our theoretical framework incorporates rich heterogeneity in firm productivity and regional fundamentals.

The rest of the paper proceeds as follows. Section 2 describes the platform, data, and measurement. Section 3 documents the empirical facts. Section 4 presents the model. Section 5 states the analytical results and defines equilibrium. Section 6 maps the evidence into the calibration. Section 7 reports the quantitative results. Section **??** compares the oligopoly baseline to the constant-markup benchmark. Section 8 concludes.

2. Background, Data, and Measurement

2.1 The Platform

The data used in this paper are collected from a professional social media platform specializing in international trade. Users registered on the platform can choose to report their firm affiliation, location (city) and industry. Registration provides access to a searchable archive of historical discussions and the ability to post new questions and replies. Many threads concern destination-specific information, trade procedures, buyer search, certification, logistics, and problem-solving around market access.

Table 1 describes the platform data. Panel A reports the distribution of discussions across 13 subforums; Panel B reports the industry distribution of platform users compared with customs exporters and the matched sample. The largest subforum is industry communication (33 percent of threads), followed by supporting services and business communication. The industry distribution closely parallels the customs data, confirming the representativeness of the platform sample.

Two features of the platform are central to the analysis. First, the discussion format is public. Even firms that do not initiate a discussion can learn from the stock of archived postings, creating positive information externalities. Second, the platform is privately operated and monetizes through advertising rather than user fees, giving it a strong incentive to maintain content quality and eliminate misinformation. In practice, the platform monitors content and organizes discussions by subforum and topic, creating a searchable information stock.

2.2 Data Sources and Sample Construction

We combine three datasets to construct our firm-level analysis sample and to discipline the quantitative exercise.

Platform data. The data cover the universe of user registrations and postings on the platform from 2003 to 2022. Our main empirical analysis focuses on 2007–2017, as posting activity is sparse before 2007. The core sample consists of 650,490 users with valid city and industry information, corresponding to 602,424 firms. Of these users, 224,718 actively post or

reply, while the remainder are passive users who may read posts but do not join discussion.

Customs data. Transaction-level Chinese customs records from 2000 to 2017. Each record reports the firm name, year, export or import value, quantity, HS product code, and destination or source country. We match firm names reported on the platform to the customs data using exact matching, fuzzy matching, and manual verification.[6] This procedure yields 89,925 matched firms, accounting for roughly 10 percent of all firms that export in the customs data over the sample window. This is a lower bound on adoption among exporters because some platform users may misspell firm names, some firms export indirectly through intermediaries, and some platform firms may not yet be direct exporters.

Input–output data. The World Input-Output Database (WIOD) and China's 2012 input–output table. These data discipline domestic supplier–buyer relationships and the sectoral structure used in the quantitative exercise. The quantitative model uses all 31 mainland Chinese provinces as domestic locations and treats the rest of the world as an additional region.

2.3 Measurement

We measure three core objects from the platform data. First, *platform adoption* is the first year in which any employee of a firm registers on the platform. Only registered users get complete access to all historical posting content and the ability to post questions.[7] Second, the *information stock* of a location-sector—a node in the platform network, where each node corresponds to one province–sector pair—is the cumulative stock of inbound responses (replies received from other nodes) available to firms in that node. Information stock increases with both active user count and posting intensity. Third, *bilateral information flows* are measured from directed posting and reply activity between source and destination location-sectors, aggregated at the region-sector-pair-period level. These bilateral flows are normalized into network shares used in the quantitative model. Trade is measured from customs data at the firm-year and province-sector-year levels.

[6]The matching procedure is as follows. First, we standardize common corporate suffixes in Mandarin. Second, we exact-match firm names, yielding 78,778 matches. Third, we fuzzy-match abbreviated names using edit distance and manual checks, yielding 9,429 additional matches. Fourth, we manually verify borderline cases using the leading firm databases Qichacha and Tianyancha, yielding 2,081 additional matches.

[7]Non-registered guest users can read only the first page of postings of recent threads.

Table 1: Platform Data: Discussion Topics and Industry Distribution

Panel A: Discussion Topics

Topics (Subforum Names)	No. Threads	Share (%)	No. Postings	Share (%)
Industry Communication	866,243	32.69	4,256,779	16.66
Supporting Services	402,294	15.18	6,490,950	25.41
Business Communication	339,732	12.82	3,903,624	15.28
E-Commerce	221,175	8.35	2,109,699	8.26
Business Opportunities	141,634	5.34	1,020,601	4.00
Import	125,254	4.73	803,929	3.15
Regional Markets	119,229	4.50	1,482,254	5.80
Recommendations This Month	117,504	4.43	1,363,103	5.34
Wonderful Life in Foreign Trade	99,229	3.74	1,608,368	6.30
Same-City Communication	84,618	3.19	1,212,354	4.75
Related Skills	73,709	2.78	618,203	2.42
Manager Communication	39,479	1.49	403,566	1.58
Foreign Trade Related	19,884	0.75	272,327	1.07
Total	2,649,984	100	25,545,756	100

Panel B: Industry Distribution

Industry (ISIC4)	Description	Platform Firm		Customs Firm		Matched Firm	
		No.	%	No.	%	No.	%
I. Production Industries							
C13-C15	Textiles, Apparel and Leather	106,080	17.61	179,351	18.84	14,996	16.68
C27-C28	Electrical Equipment and Machinery	94,398	15.67	147,241	15.47	14,617	16.25
C24-C25	Metal Products	64,960	10.78	79,449	8.35	9,157	10.18
C26	Computer, Electronic and Optical	41,832	6.94	140,760	14.79	15,152	16.85
C22-C23	Rubber, Plastic, Mineral	39,625	6.58	93,035	9.77	9,206	10.24
C31-C32	Furniture and Other Mfg.	38,133	6.33	72,878	7.66	7,633	8.49
C19-C21	Petroleum, Chemicals, Pharma.	36,782	6.11	68,936	7.24	7,762	8.63
C10-C12	Food Products	18,647	3.10	44,452	4.67	2,518	2.80
C16-C17	Wood and Paper Products	17,678	2.93	40,816	4.29	2,991	3.33
C29-C30	Vehicles and Transport Equip.	8,911	1.48	37,134	3.90	2,961	3.29
A01, A03	Agriculture, Fishing	8,485	1.41	22,610	2.38	1,346	1.50
B	Mining and Quarrying	8,200	1.36	11,143	1.17	774	0.86
Others	Electricity, Support, Others	6,783	1.13	14,133	1.48	812	0.90
II. Service Industries							
H49-H53	Transportation	51,481	8.55				
G45-G47	Wholesale and Retail Trade	40,918	6.79				
M74-M75	Professional and Technical	8,168	1.36				
J62-J63	Programming and Information	7,361	1.22				
III. Other Services							
—	Legal, Finance, Other Services	3,982	0.66				
Total		602,424	100.00	951,938	100.00	89,925	100.00

Notes: Panel A summarizes the number of threads and postings by subforum, ranked by popularity. Panel B presents the firm count in the platform, the customs data (2000–2017), and the matched sample. Platform firm industry is determined by the most frequently reported industry by its workers (frequency method). Customs and matched firm industry are determined by the firm's largest HS-2 digit industry by trade value. The correlation between the two methods is 0.98. "Others" include ISIC D35, J58, N, P85, R_S and HS 97–98. "Other services" groups sectors without a matched IO code in the 2006 WIOD 16-sector concordance (Legal, Finance, and residual service categories).

3. Empirical Facts

We organize the empirical evidence in four steps. We first estimate the causal effect of platform adoption on exports (Fact 1), then document who adopts the platform and where (Fact 2), and characterize the structure of information flows—their geographic decay (Fact 3) and alignment with production networks (Fact 4).

Fact 1: Platform Adoption and Exports Growth

Empirical strategy. We estimate the causal effect of platform adoption on firm-level export outcomes using a stacked cohort event-study design with propensity-score matching. The design exploits the staggered timing of firms' first platform registration across cohorts and sectors. We define firm i's treatment year as the first calendar year in which any employee registers on the platform, imposing a monotone-treatment (absorbing-state) assumption.

The stacked approach follows Cengiz et al. (2019), who construct clean two-by-two comparisons for each treatment event and stack them. Our specification uses an asymmetric $[-3, +4]$ event window, with three pre-treatment leads to test parallel trends and four post-treatment periods to trace the dynamic evolution of effects. We require both treated and control firms to be observed continuously throughout the event window plus a one-year margin on each side, for a total of 10 consecutive years of positive exports.[8] Within each cohort–sector cell, we estimate a logit propensity score using five pre-treatment covariates measured at $k = -1$: log export value, export growth (change in log exports over the pre-period), log destination count, log product count, and log average unit price. Each treated firm is matched to its nearest-neighbor control on the estimated propensity score.

On the stacked matched panel, we estimate:

$$y_{sit} = \sum_{\substack{k=-3 \\ k\neq -1}}^{4} \beta_k \cdot \mathbf{1}\{t - c_s = k\} \cdot T_{si} + \alpha_{si} + \delta_{st} + \varepsilon_{sit}, \tag{3.1}$$

[8]The 10-year continuity requirement ($[-3, +4]$ event window = 8 years, plus a one-year margin on each side for computing lagged growth rates) eliminates compositional changes across the event window: every firm contributes to every event-time coefficient. It also guarantees well-defined log outcomes in every period, since every firm has strictly positive exports, destinations, and products throughout the window, eliminating the log-zero concerns raised by Chen and Roth (2024). Of the 4,781 candidate treated firms that satisfy this requirement, 4,764 are successfully matched.

where y_{sit} is the outcome for firm i in matched set s at calendar year t; $T_{si} \in \{0, 1\}$ is a treatment indicator equal to one if firm i is the treated member of matched pair s; c_s denotes the treatment year for cohort s, so $\mathbf{1}\{t - c_s = k\}$ selects event time k relative to treatment; β_k is the event-time coefficient tracing the dynamic treatment effect; α_{si} are firm-by-matched-set fixed effects absorbing time-invariant differences within each pair; and δ_{st} are matched-set-by-calendar-year fixed effects absorbing common time shocks within each cohort–sector cell. The omitted reference period is $k = -1$. Standard errors are two-way clustered by matched-set and firm. We estimate equation (3.1) on six outcomes: log total export value, log destination count, log HS product count, log average unit price, and two intensive margins (log export per destination and log export per product). Appendix Table A1 compares the matched sample with the full customs population.

Results: Export value. Table 2 reports the baseline estimates for log export value. The design yields 4,764 treated firms matched to 4,764 controls, for a total of 76,224 firm–year observations. The identifying assumption is that, conditional on matching covariates and fixed effects, treated and control firms would have followed parallel export trajectories absent treatment. The pre-treatment coefficients at $k = -3$ and $k = -2$ are close to zero and precisely estimated, and the confidence intervals comfortably include the origin. The dynamic pattern before treatment is flat, consistent with the parallel-trends assumption.

The post-treatment path tells a clear story. In the year of first registration, exports rise by about 6 percent relative to the matched control group. The effect then builds steadily over the next three years, reaching roughly 10 percent after one year, 12 percent after two years, and stabilizing near 11–13 percent at longer horizons. Two features of this trajectory are informative about the underlying mechanism. First, the gradual build-up is inconsistent with a one-time visibility or demand shock, which would produce an immediate level shift that does not grow. Instead, the rising profile over three to four years points to a learning channel: firms progressively accumulate destination-specific knowledge from the platform's archived discussions and translate that knowledge into expanded export activity over time. Second, the slight moderation at $k = 4$ suggests diminishing marginal returns to platform information as firms exhaust the most accessible knowledge gains, though the estimate remains well within the confidence band of the peak.

To place the magnitude in context, our estimates are comparable to causal evaluations of formal export promotion programs. Munch and Schaur (2018) find that Danish trade promotion raises medium-sized firm exports by 12–16 percent after two years. Carballo et al. (2022) estimate a 17 percent increase from joining ConnectAmericas. Broocks and Van Biesebroeck (2017) find an 8.5 percentage point increase in the probability of entering a new export market from Belgian trade promotion. The mean post-treatment effect in our baseline—about 10 percent—is of a similar order of magnitude. The fact that voluntary adoption of a digital platform produces export gains comparable to government-sponsored interventions is itself an economically significant finding.

Decomposition: Extensive vs. intensive margins. Panel A of Table 2 reports total exports alongside the two extensive margins; Panel B reports prices and the two intensive margins. All six columns are estimated on the identical matched sample. Figure 1 plots the corresponding event-study paths.

The destination margin dominates. Firms that join the platform serve about 7 percent more foreign destinations in the treatment year—roughly one-half additional market for a firm that previously exported to five—and the gap widens to 13 percent within two years. The effect is persistent and, if anything, still growing at $k = 4$. Product scope expands as well, though more modestly: about 3 percent on impact, rising to 8 percent after two years.

The intensive margins are flat or small. Log export per destination (column 5) is indistinguishable from zero at every horizon, with a mean post-treatment coefficient of -0.014. Log export per product (column 6) shows a modest positive response of 0.035, driven by gains at $k = 2$ and $k = 3$. Unit prices (column 4) show no discernible response: the mean post-treatment coefficient is -0.015, and the confidence intervals are centered on zero throughout.

The log-additive identity $\ln(\text{export}) = \ln(\text{destinations}) + \ln(\text{export/destinations})$ allows a precise decomposition. The destination-extensive margin accounts for more than the entire export increase: the mean post-treatment gain in log destinations (0.116) exceeds the mean gain in total exports (0.103), while the destination-intensive margin is -0.014. Along the product dimension, the extensive margin accounts for roughly two-thirds of the total effect (0.068) and the intensive margin contributes the remaining third (0.035).

Table 2: Effect of Platform Adoption on Firm Exports

Panel A: Total and extensive margins

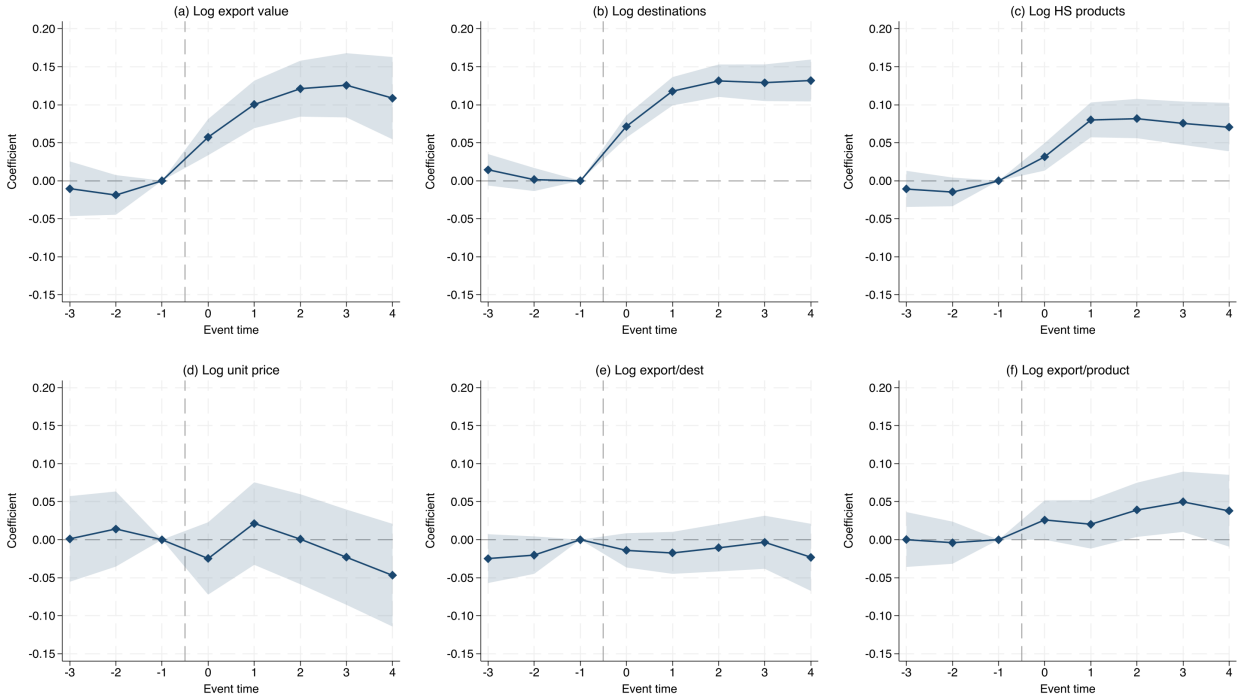
	(1) Export value		(2) Destinations		(3) Products (HS)	
Event time	$\hat{\beta}$	SE	$\hat{\beta}$	SE	$\hat{\beta}$	SE
$t-3$	−0.011	(0.018)	0.014	(0.011)	−0.011	(0.012)
$t-2$	−0.019	(0.013)	0.002	(0.008)	−0.015	(0.010)
t	0.057***	(0.012)	0.071***	(0.007)	0.031***	(0.009)
$t+1$	0.100***	(0.016)	0.118***	(0.010)	0.080***	(0.012)
$t+2$	0.121***	(0.019)	0.132***	(0.011)	0.082***	(0.013)
$t+3$	0.126***	(0.022)	0.129***	(0.012)	0.076***	(0.015)
$t+4$	0.109***	(0.028)	0.132***	(0.014)	0.071***	(0.016)
Joint pre p	0.364		0.316		0.311	
Mean post	0.103		0.116		0.068	

Panel B: Price and intensive margins

	(4) Unit price		(5) Export/dest		(6) Export/product	
Event time	$\hat{\beta}$	SE	$\hat{\beta}$	SE	$\hat{\beta}$	SE
$t-3$	0.001	(0.029)	−0.025	(0.016)	0.000	(0.018)
$t-2$	0.014	(0.025)	−0.020	(0.013)	−0.004	(0.014)
t	−0.025	(0.024)	−0.014	(0.011)	0.026*	(0.013)
$t+1$	0.021	(0.028)	−0.017	(0.014)	0.020	(0.016)
$t+2$	0.001	(0.030)	−0.011	(0.016)	0.039**	(0.018)
$t+3$	−0.023	(0.032)	−0.003	(0.018)	0.050**	(0.020)
$t+4$	−0.047	(0.034)	−0.023	(0.023)	0.038	(0.024)
Joint pre p	0.821		0.200		0.942	
Mean post	−0.015		−0.014		0.035	
Treated pairs	4,764					
Observations	76,224					
Window	[−3, +4], support-only, 1-NN pscore matching					

Notes: Stacked cohort event-study with propensity-score matching. All six outcomes estimated on the identical matched sample. (1) log total export value. (2) log number of destination countries. (3) log number of HS products. (4) log average unit price. (5) log export per destination = (1)−(2). (6) log export per product = (1)−(3). Omitted period: $k = -1$. Propensity scores estimated within cohort–sector cells using pre-treatment export level, growth, destination count, product count, and unit price. Both groups trimmed to common support. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. SE two-way clustered by matched-set and firm.

Figure 1: Effect of Platform Adoption on Export Margins



Notes: All six panels share the same vertical axis. Panel A (top row): (a) log total export value, (b) log number of destinations, (c) log number of HS products. Panel B (bottom row): (d) log average unit price, (e) log export per destination, (f) log export per product. Coefficients in (b)+(e) and (c)+(f) each sum to (a) by construction. All panels use the same matched sample (N = 76,224; 4,764 pairs). Shaded areas: 95% CI. Two-way clustered SE by matching set and firm.

Interpretation. The combination of large extensive-margin effects, null prices, and flat intensive margins points to a specific mechanism: the platform reduces the *fixed costs* of discovering and entering new markets—search frictions, regulatory uncertainty, destination-specific know-how—rather than lowering variable trade costs or matching firms to higher-paying buyers. In the Chaney (2008) framework, lower fixed entry costs expand the set of profitable destinations without affecting per-unit prices, which is exactly the pattern we observe. The null price response rules out a buyer-matching channel (which would raise unit values) and a variable-cost channel (which would lower them). The flat destination-intensive margin rules out explanations in which the platform helps firms sell more to existing customers; instead, it helps firms find *new* customers in *new* markets. This is also consistent with the export promotion literature: Volpe Martincus and Carballo (2008) and Broocks and Van Biesebroeck (2017) find that trade promotion programs expand the extensive margin with little or no intensive-margin effect. The formal model in Section 4 makes this mechanism precise and shows how the empirical patterns discipline the model's key parameters.

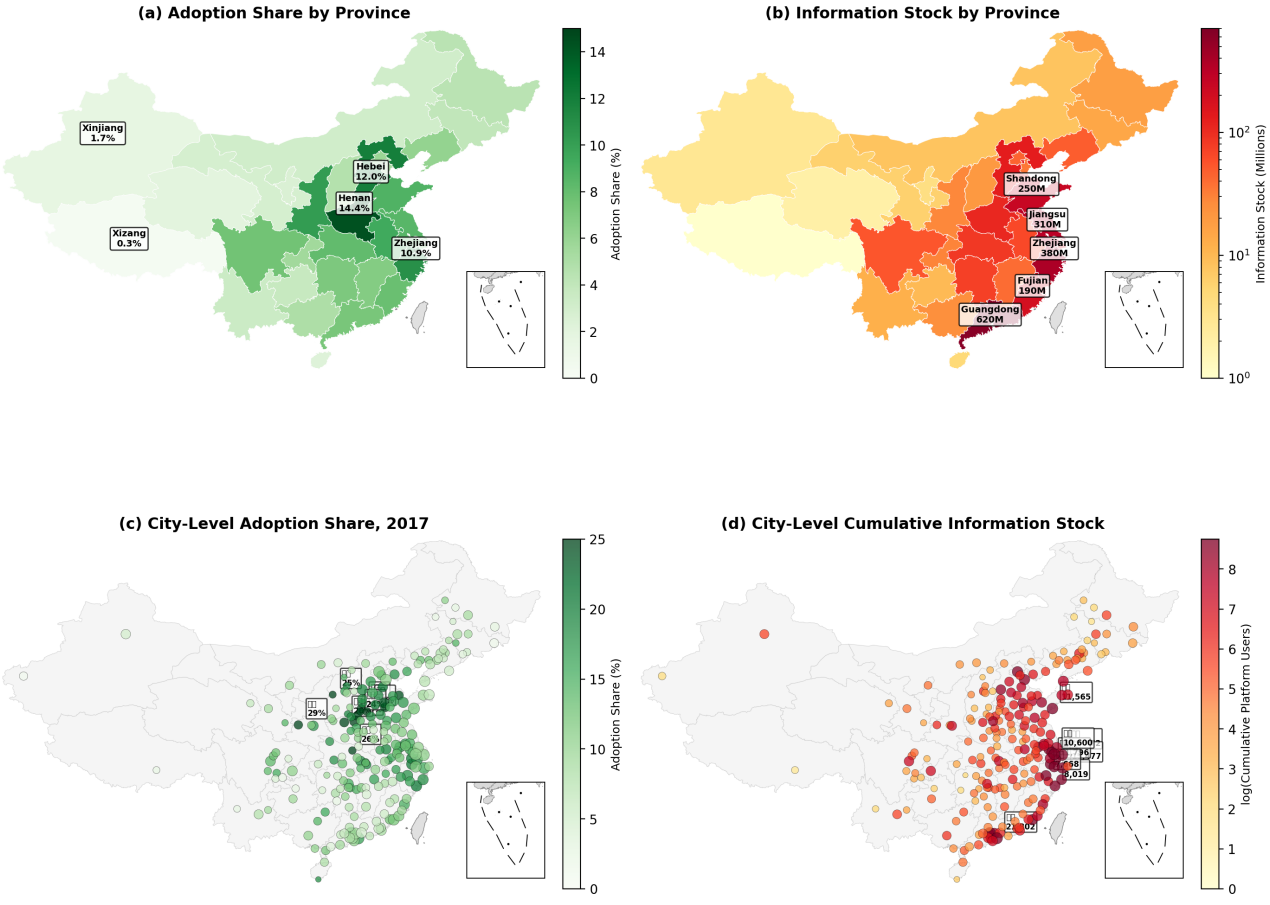
Fact 2: Adoption Is Uneven Across Regions and Firms

Regional adoption patterns. Since the development of the internet, people have debated whether it equalizes access to information (Rosenblat and Mobius, 2004). Figure 3 shows that the share of exporting firms that have at least one employee on the platform increases significantly with city-level GDP per capita and internet broadband access. Though registering on the platform is free, effective adoption requires both internet infrastructure and awareness of the platform's benefits. This pattern is somewhat surprising: one might expect an online platform to disproportionately benefit low-income regions where firms have less offline information. Two forces work against equalization. First, less developed regions have lower internet access and digital literacy, raising the effective cost of adoption. Second, firms in less developed regions are less likely to receive help from peers because the expected economic benefit from helping them is lower—an implication of the endogenous knowledge-posting mechanism in the model.

Figure 2 maps these adoption patterns geographically. Panels (a) and (b) display province-level adoption shares and information stocks, respectively. Henan (14.4 percent), Hebei (12.0 percent), and Zhejiang (10.9 percent) lead in adoption, while Xizang (0.3 percent) and Xin-

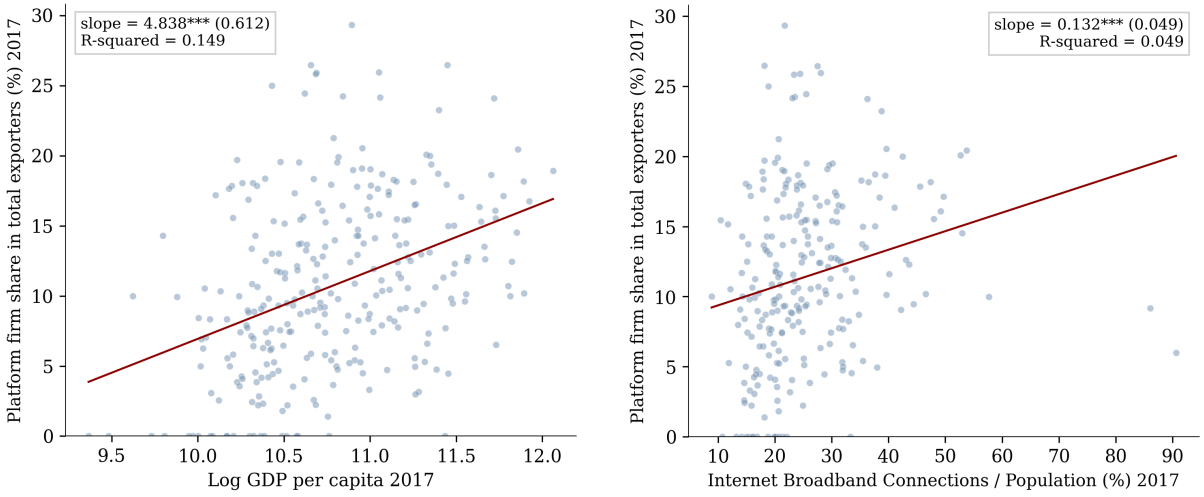
jiang (1.7 percent) lag far behind. Information stocks are highly concentrated: Guangdong alone accounts for 620 million reply-equivalents, and the top five provinces hold 85.3 percent of the national total. Panels (c) and (d) disaggregate to the prefecture-city level for 2017. The city-level maps reveal substantial within-province heterogeneity that the province aggregates conceal. Coastal cities cluster in the upper right of both distributions, while interior cities—even within high-adoption provinces—exhibit markedly lower adoption and thinner information stocks. The east–west gradient in panel (d) is consistent with a distance-to-coast effect: cities farther from major ports accumulate fewer platform interactions, likely because fewer local exporters generate the knowledge externalities that attract subsequent users.

Figure 2: Geography of Platform Adoption and Information



Notes: Panels (a)–(b): province-level measures, pooled 2007–2017. Panel (a) maps the adoption share (percent of exporting firms with at least one registered platform user); darker green indicates higher adoption. Panel (b) maps cumulative information stocks (total inbound reply-equivalents) on a logarithmic color scale; top 5 provinces are labeled. Panels (c)–(d): prefecture-city level, 2017. Panel (c) maps city-level adoption shares; bubble size proportional to the number of exporting firms in each city. Panel (d) maps cumulative platform users (a proxy for the local information stock) on a log scale. Top 8 cities by each measure are labeled. Data sources: customs–platform matched records and platform posting data. The boundaries shown do not imply any judgment on the legal status of any territory.

Figure 3: Regional Adoption Correlates



Notes: City-level platform adoption share among exporters in 2017 vs. log GDP per capita (left, slope = 4.838, $p < 0.01$) and internet broadband connections per capita (right, slope = 0.132, $p < 0.01$).

Appendix Figure A1 traces the geographic diffusion at the prefecture-city level from 2009 to 2017: adoption begins in a handful of coastal hubs and spreads inland, though the east–west intensity gradient remains pronounced throughout.

Firm-level selection into adoption. A natural question is whether firm size predicts adoption. If joining the platform requires an upfront investment of time and effort—learning to navigate discussions, building a reputation, monitoring threads—then firms with larger export portfolios have more to gain from that investment, making them more likely to adopt. We test this using a cross-sectional logit regression of the adoption indicator on pre-treatment firm characteristics:

$$\Pr(\text{Adopt}_{it} = 1) = \Lambda(\alpha + \beta \ln X_{i,t-1} + \gamma' Z_{i,t-1} + \delta_s + \varepsilon_{it}) \tag{3.2}$$

where $\Lambda(\cdot)$ is the logistic CDF, $X_{i,t-1}$ is pre-treatment export value, $Z_{i,t-1}$ are additional firm controls, and δ_s are sector fixed effects (included in columns 4–5 of Table 4). Table 3 compares adopters and non-adopters on pre-treatment observables. Across 193,503 firm-cohort observations (7,343 treated, 186,160 control), adopters have 1.56 higher log exports ($t = 56.3$), 1.17 more log destinations ($t = 74.9$), 0.66 more log products ($t = 38.5$), and 0.39 higher log unit prices ($t = 10.7$). All differences are highly significant.

Table 3: Pre-Treatment Firm Characteristics: Treated vs. Control

	Treated (Platform)		Control		Difference	
	Mean	SD	Mean	SD	Δ	t-stat
Log export value	14.98	1.79	13.42	2.14	1.56	56.3
Log destinations	2.53	1.02	1.37	1.14	1.17	74.9
Log products (HS)	2.09	1.31	1.43	1.22	0.66	38.5
Log unit price	3.27	2.70	2.89	2.65	0.39	10.7
N firms	7,343		186,160			

Notes: Pre-treatment characteristics measured in the year before each cohort's adoption year. Treated firms join the platform in their cohort year; controls are never-treated exporters satisfying event-study eligibility criteria. Each firm appears once (first cohort appearance). t-statistics from two-sample t-tests with unequal variances.

Table 4 reports logit regressions of the adoption indicator on these pre-treatment characteristics. Column (1) uses log exports alone: a one-unit increase in log exports raises the log-odds of adoption by 5.6. Columns (2)–(4) progressively add destination scope, product scope, unit prices, export growth, and sector fixed effects. In the most demanding specification—column (5) with sector $\times$ cohort fixed effects—log exports, log destinations, and export growth are all strongly predictive: a one-unit increase in log exports raises the log-odds by 3.7, and the pseudo R^2 jumps to 0.28. The average marginal effects (AME column) indicate that a one-unit increase in log exports raises the adoption probability by 14.8 percentage points; a one-unit increase in log destinations adds 6.6 percentage points.

Table 5 shows the pattern non-parametrically. Sorting firms into quintiles by pre-treatment export value, the adoption rate rises monotonically from 0.7% in the bottom quintile to 8.4% in the top quintile—a 12-fold gradient. This monotone selection pattern is consistent with the model's cutoff structure (Proposition 5.2): in equilibrium, only firms above an endogenous productivity threshold find it profitable to pay the fixed learning cost.

Table 4: Selection into Platform Adoption: Logit Regressions

Dep. var.: 1(adopt by year t)	(1)	(2)	(3)	(4)	(5)
Log export	5.562***	2.287***	2.241***	2.229***	3.740***
	(0.092)	(0.128)	(0.127)	(0.128)	(0.142)
Log destinations		1.454***	1.469***	1.474***	1.659***
		(0.041)	(0.041)	(0.042)	(0.045)
Log products		-0.113***	-0.163***	-0.164***	-0.040
		(0.024)	(0.025)	(0.025)	(0.028)
Log unit price			0.040***	0.047***	0.030***
			(0.005)	(0.006)	(0.006)
Export growth			0.032***	0.033***	0.228***
			(0.009)	(0.009)	(0.010)
Sector FE	No	No	No	Yes	No
Sector × Cohort FE	No	No	No	No	Yes
Observations	193,500	123,187	123,187	123,187	123,187
Pseudo R^2	0.066	0.079	0.081	0.082	0.281

Notes: Logit coefficients. Dependent variable equals one if firm joins the platform platform. Pre-treatment characteristics measured in the year before each cohort's adoption year. Each firm appears once (first cohort). Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: platform Adoption Rate by Pre-Treatment Export Quintile

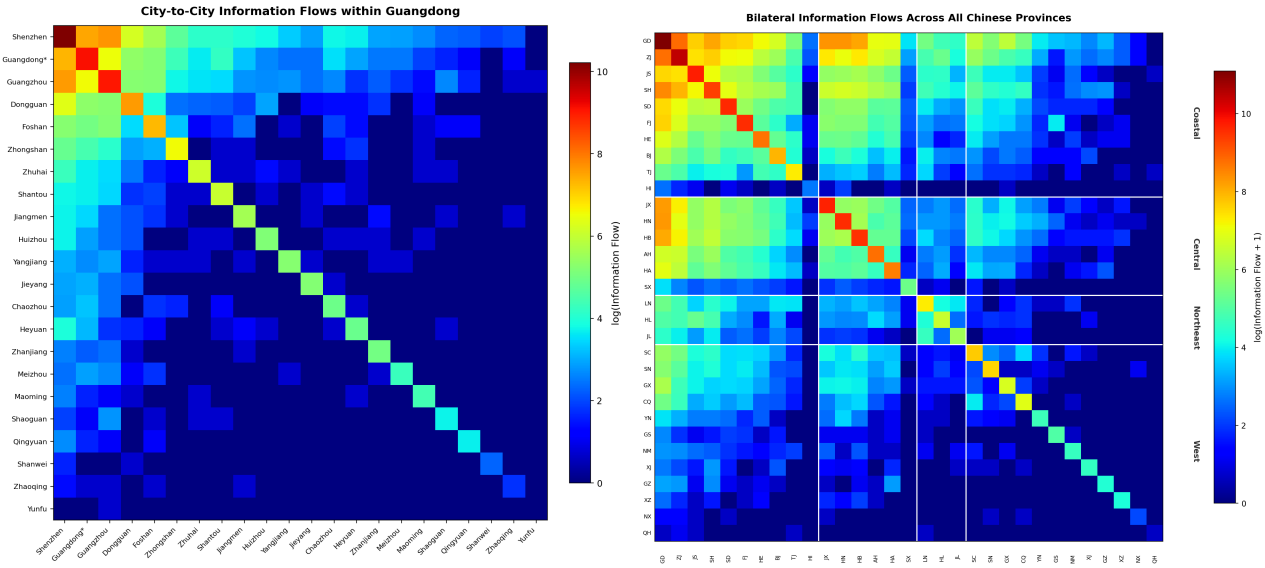
	Q1 (smallest)	Q2	Q3	Q4	Q5 (largest)
Adoption rate (%)	0.7	1.7	3.3	4.9	8.4
N firms	38,701	38,701	38,700	38,703	38,698

Notes: Share of firms that join the platform platform, by quintile of pre-treatment export value. Quintiles based on export value in the year before each cohort's adoption year. Each firm appears once (first cohort appearance). The adoption rate rises monotonically from 0.7% (Q1) to 8.4% (Q5), a 12-fold gradient.

Fact 3: Information Flows Are Geographically Local

Geography is a key barrier to both information and trade flows. We measure information flows at the sector-region-pair-period level by aggregating postings. Figure 4 shows the bilateral information flow heatmap. The concentrated diagonals show that within-city and within-province flows dominate, and the Coastal block is far more active than interior regions. Appendix Figure A2 provides a complementary network visualization.

Figure 4: Bilateral Information Flows



(a) Within Guangdong: City-to-City (b) All 31 Provinces

Notes: Panel (a): city-to-city information flows within Guangdong, the province with the largest information stock. Cities are ordered by total flow volume; Shenzhen, Guangzhou, and Dongguan dominate both as sources and destinations. The strong diagonal reflects within-city information concentration; off-diagonal flows reveal the Pearl River Delta cluster (Shenzhen–Guangzhou–Dongguan–Foshan). Panel (b): province-to-province bilateral flows across all 31 provinces, grouped by macro-region (Coastal, Central, Northeast, West). White lines delineate macro-region boundaries. Colors encode log(flow count + 1). The Coastal block (upper left) dominates; cross-region flows are asymmetric, with Coastal provinces sending more information to Central and West provinces than they receive. Data source: platform posting and reply records, 2007–2017.

To formally test the local diffusion of information, we estimate the following gravity equation using Poisson pseudo-maximum-likelihood (PPML):

$$\text{InfoFlow}_{aibjt} = \exp\left(\beta^{Dist} \log \text{Dist}_{ij} + Z_{ij} + \delta_{abt} + \xi_{ait} + \eta_{bjt}\right) \times \epsilon_{aibjt} \tag{3.3}$$

where InfoFlow_{aibjt} is the total responsive posting number from location i sector a to location j sector b in period t, Dist_{ij} is the geographic distance in kilometers, Z_{ij} includes same-city

and same-province dummies, and the fixed effects absorb sector-pair-time, origin city-sector-time, and destination city-sector-time variation.

Table 6 presents the results. The elasticity of information flows with respect to distance is −0.064 (column 1), indicating that a 10 percent increase in distance reduces information flows by 0.64 percent. Column 2 weights postings by character count, yielding −0.085. Appendix Table **??** places these estimates alongside comparable measures from the literature.

Table 6: Gravity Regressions: Distance Elasticity of Information and Trade Flows

	Panel A: City-Sector Level		*Panel B: Province-Sector Level*	
	(1) Info: unwt.	(2) Info: wtd.	(3) Information	(4) Trade
Log Distance	−0.064***	−0.085***	−0.338***	−0.672***
	(0.011)	(0.018)	(0.008)	(0.002)
Same City	2.523***	2.866***		
	(0.069)	(0.100)		
Same Province	0.284***	0.322***		
	(0.040)	(0.069)		
Observations	279,172	279,172	73,902	3,335,724
Origin-Sector-Time FE	Yes	Yes	Yes	Yes
Dest.-Sector-Time FE	Yes	Yes	Yes	Yes
Sector-Pair-Time FE	Yes	Yes	Yes	Yes

Notes: PPML gravity regressions. Panel A: city-sector-pair level, dependent variable is bilateral posting count. Column (1) uses unweighted counts; column (2) weights by character length. Panel B: province-sector-pair level. Column (3): information flows; column (4): trade flows. The ratio of distance elasticities in Panel B (0.338/0.672 = 0.503) disciplines the curvature of information provision in the model. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. SE clustered at pair level.

The comparison between Panels A and B, and between columns (3) and (4), reveals two key patterns. First, information flows decay with distance at roughly half the rate of trade flows (0.338 vs. 0.672), consistent with information being partially but not fully digitizable. Second, the city-level regressions in Panel A show that crossing a city boundary reduces information flows by a factor of $\exp(-2.52) \approx 12$, far exceeding the continuous distance effect. The ratio of information-to-trade distance elasticities in Panel B disciplines the curvature of information provision in the model.

Viewed against the broader literature (Appendix Table **??**), three comparisons stand out. First, *information flows are substantially less distance-sensitive than trade.* The mean trade elasticity across 1,467 estimates is −0.9 (Disdier and Head, 2008), and −1.1 on average across 2,508 estimates (Head and Mayer, 2014); our province-level information elasticity of −0.338 is less

than half this magnitude. Even within China—where domestic trade frictions are lower—Arkolakis, Huneeus and Miyauchi (2023) estimate −1.334 at the municipality level, roughly four times our city-level information estimate. This gap is consistent with digital communication reducing the cost of transmitting knowledge across space, even while platform users remain concentrated locally. Second, *our information elasticity is larger than R&D-based proxies but smaller than financial flows.* Keller (2002) estimates a distance elasticity of −0.07 for bilateral TFP growth with respect to distance-weighted foreign R&D—one order of magnitude below our platform estimate—reflecting that technology embodied in traded goods diffuses broadly regardless of origin. By contrast, Portes and Rey (2005) find a distance elasticity of −0.79 for cross-border equity flows, which combine information frictions with formal institutional barriers; our measure, restricted to voluntary digital communication, falls well below this. Third, *the information-to-trade elasticity ratio of* 0.503 *is economically meaningful.* If information were fully frictionless (pure digital good), the ratio would approach zero; if it behaved like a physical good, the ratio would approach one. The intermediate value of 0.503 confirms that geography still matters for knowledge—consistent with Jaffe, Trajtenberg and Henderson's (1993) original finding that patent citations are spatially concentrated—but that digital platforms do compress the distance friction substantially relative to trade.

This finding motivates our theory that firms have stronger incentives to share information with geographically closer peers due to higher expected business returns. An input supplier has greater incentive to assist a nearby input buyer in expanding exports because the buyer's success increases future demand for the supplier's inputs. Lower geographic distance enhances incentive compatibility among firms within supply chains.

Two extensions in Appendix **??** deepen this finding. First, the decay profile is sharply nonlinear: replacing log distance with distance-bin dummies reveals that 88 percent of total attenuation occurs at the city boundary, with little additional decay over the next 900 km—geography operates primarily as a discrete barrier, not a continuous friction (Appendix Figure A3). Second, expert knowledge travels farther: the distance elasticity for confirmed knowledge replies (−0.24) is 43 percent smaller than the aggregate, consistent with buyers crossing geographic boundaries to access scarce expertise.

External economies of scale. The geographic concentration of information flows has a direct consequence for returns to platform adoption. If the platform generates positive externalities—so that one firm's postings benefit others in the same location-sector—then returns to adoption should increase in the density of the local information stock, and those returns should be largest where information is richest. This is the mechanism emphasized in the external economies of scale literature (Bartelme et al., 2025; Lashkaripour and Lugovskyy, 2023).

We test this prediction using a firm-level panel regression that interacts adoption status with the local information stock:

$$\ln Y_{f,t} = \beta_1 \, Adopted_{f,t} + \beta_2 \, Adopted_{f,t} \times \ln S_{r(f),s(f),t} + \alpha_f + \delta_{r,t} + \varepsilon_{f,t}, \tag{3.4}$$

where $Y_{f,t}$ is firm f's total exports, $Adopted_{f,t}$ is a dummy equal to one once any employee of firm f has registered on the platform, $S_{r(f),s(f),t}$ is the cumulative information stock (total inbound replies) in firm f's province-sector cell as of year t, α_f are firm fixed effects, and $\delta_{r,t}$ are province-year fixed effects. The coefficient β_2 tests whether the returns to adoption are increasing in the local information stock. If $\beta_2 > 0$, the platform exhibits external economies of scale: adoption is more valuable where more knowledge has been collectively deposited.

Table 7 reports the results (N = 2,791,754 firm-year observations). Column (1) shows the unconditional adoption effect ($\hat{\beta}_1 = 0.326^{***}$). Column (2) introduces the interaction: the coefficient on adoption alone falls close to zero ($\hat{\beta}_1 = 0.033$, n.s.), while the interaction term is $\hat{\beta}_2 = 0.031^{***}$ (SE = 0.005)—a 10 percent increase in local information stock is associated with a 0.31 percentage point larger adoption premium. The returns to adoption are negligible in locations with thin information stocks and large in information-rich environments, exactly the pattern expected under positive externalities.

A natural concern is that the information-stock variable may proxy for the overall density of local export activity rather than platform-specific knowledge. Locations with many exporters mechanically have more platform users, more postings, and more information stock. To rule out this agglomeration confound, columns (3) and (4) add the interaction of adoption with the log number of exporters in the same province-sector-year cell. This control absorbs any returns to adoption that arise from being located in a thick export cluster rather than

from the platform's information content per se. The key finding is that the information-stock interaction remains stable and precisely estimated after controlling for local exporter density. In column (3), the information-stock interaction is 0.029 (compared with 0.031 without the control), while the exporter-density interaction is small and not statistically distinguishable from zero. Column (4) uses weighted information stocks and shows the same pattern. This indicates that the returns to adoption are driven by the quality and quantity of platform-specific knowledge, not by the general thickness of the local export environment. Column (5) weights postings by character length without the exporter-density control for comparison.

The interaction coefficient of 0.029–0.031 implies highly concave returns to information: the marginal value of additional platform content diminishes rapidly with stock size. Section 6 uses this estimate to discipline the information technology in the model.

Table 7: Impact of Information Stock on the Returns to Platform Adoption

Dependent variable: log firm exports	(1)	(2)	(3)	(4)	(5)
Adoption	0.326***	0.033	0.018	−0.082	−0.063
	(0.008)	(0.049)	(0.051)	(0.060)	(0.058)
Adoption × ln(Info Stocks)		0.031***	0.029***		
		(0.005)	(0.006)		
Adoption × ln(Wtd Info Stocks)				0.024***	0.026***
				(0.005)	(0.004)
Adoption × ln(Local Exporters)			0.008	0.009	
			(0.012)	(0.012)	
Observations	2,791,754	2,791,754	2,791,754	2,791,754	2,791,754
Adjusted R^2	0.691	0.691	0.691	0.691	0.691
Province-Year FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable: log firm exports. "Info Stocks" is the cumulative stock of responses to postings in the firm's province-sector-year cell. "Wtd Info Stocks" weights postings by character length. "Local Exporters" is the number of active exporters in the same province-sector-year cell from customs data. Columns (3)–(4) control for the interaction of adoption with log local exporter count to separate the platform information channel from general agglomeration effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors in parentheses.

Fact 4: Knowledge Transmission Concentrates Along Production Networks

We examine how value-chain relationships shape sector-to-sector information flows on the platform. For each directed sector pair (i, j) among 22 WIOD goods sectors, we construct

three measures using the platform data and the 2006 WIOD input-output table:

$$\text{InfoNet}_{ij} = \frac{\text{InfoFlow}_{ij}}{\text{InfoSupply}_i}, \quad \text{Down}_{ij} = \frac{\text{InputFlow}_{ij}}{\text{InputSupply}_i}, \quad \text{Up}_{ij} = \frac{\text{InputFlow}_{ji}}{\text{InputUsage}_i}.$$

InfoNet_{ij} is the share of sector i's outbound knowledge-sharing posts directed to sector j (mean = 0.043, s.d. = 0.052). Down_{ij} measures how much of i's output is purchased by j (downstream linkage). Up_{ij} measures how much of i's inputs come from j (upstream linkage).

OLS specification. The baseline test regresses information network shares on production-network linkages. We estimate two separate regressions:

$$\text{InfoNet}_{ij} = \alpha_D + \beta_D \, \text{Down}_{ij} + \varepsilon_{ij}^D, \tag{3.5}$$

$$\text{InfoNet}_{ij} = \alpha_U + \beta_U \, \text{Up}_{ij} + \varepsilon_{ij}^U, \tag{3.6}$$

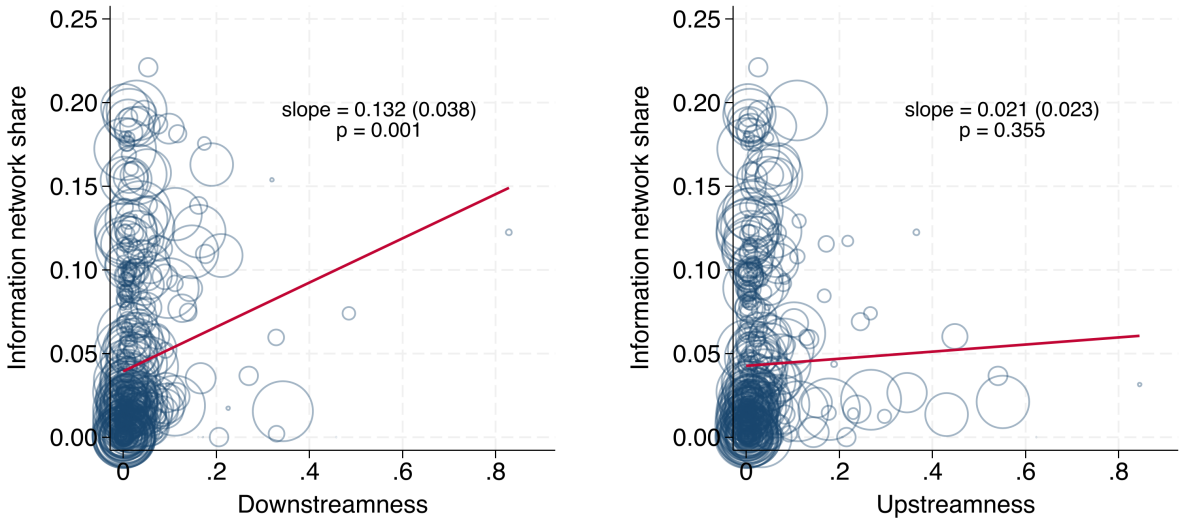
and a joint specification with both regressors. The unit of observation is a directed sector pair (i, j) with $i \neq j$, yielding $N = 22 \times 21 = 462$ observations. Standard errors are heteroskedasticity-robust (HC1). The key prediction is $\beta_D > 0$ (suppliers teach downstream buyers) and $\beta_U = 0$ (no reverse flow from input suppliers to the purchasing sector).

Figure 5 plots the raw correlations. Information sharing is significantly and positively correlated with downstream relationships ($\hat{\beta}_D = 0.132$, SE = 0.038, $p = 0.001$), while there is no significant correlation with upstream relationships ($\hat{\beta}_U = 0.021$, SE = 0.023, $p = 0.355$). This downstream–upstream asymmetry rules out passive learning—where buyers learn by examining the intermediate products they purchase—and instead supports an *intentional-teaching* mechanism: suppliers actively share export knowledge with downstream buyers because the buyer's export success generates future demand for the supplier's inputs. Hajdini et al. (2025) provide experimental evidence from New Zealand customer-supplier pairs that direct inter-firm communication—not observable actions—is the primary mechanism through which information diffuses along supply chains, consistent with our finding that intentional posting drives the downstream flow.

The pattern is economically large. A one-standard-deviation increase in downstreamness (0.067) raises the predicted information network share by 0.9 percentage points—roughly

one-fifth of the mean share (4.3%). The asymmetry is also statistically sharp: in a joint regression, a Wald test rejects $\beta_D = \beta_U$ ($F = 6.02$, $p = 0.014$).

Figure 5: Production-Network Information Flows



Notes: Each dot is a sector pair ($N = 462$, 22 WIOD goods sectors × 21 counterpart sectors). Circle size is proportional to the teaching sector's total information supply (cumulative reply count on the platform). Left: InfoNet_{ij} vs. downstreamness (slope = 0.132, $p = 0.001$). Right: InfoNet_{ij} vs. upstreamness (slope = 0.021, $p = 0.355$). Red line: unweighted OLS fit. Robust SE.

Sector-pair examples. The downstream teaching pattern is visible in specific sector pairs. The strongest information flows are directed *into* textiles and apparel (WIOD sector 6)—China's largest export category. Paper suppliers (sector 8), electronics firms (sector 17), and rubber/plastics manufacturers (sector 13) all direct 19–22% of their outbound knowledge to the textile sector, consistent with upstream suppliers sharing packaging, labeling, and compliance knowledge with downstream apparel exporters. At the other end, the strongest downstream input linkage—fishing (sector 3) to food processing (sector 5), where 83% of fishing output goes to food—is accompanied by a 12% information network share, confirming that production-chain proximity predicts knowledge flows even outside manufacturing.

Instrumental variable estimation. A concern with the OLS estimates is that unobserved common factors—such as shared export destinations or correlated demand shocks—could drive both input-output linkages and information flows. We address this using shift-share

instruments in the spirit of Bartik (1991). The instruments for Down_{ij} and Up_{ij} are:

$$Z_{ij}^{\text{Down}} = \text{Down}_{ij}^{2001} \times \Delta \ln D_j^{\text{foreign}}, \tag{3.7}$$

$$Z_{ij}^{\text{Up}} = \text{Up}_{ij}^{2001} \times \Delta \ln D_j^{\text{foreign}}, \tag{3.8}$$

where Down_{ij}^{2001} and Up_{ij}^{2001} are the pre-WTO input-output shares from the 2001 WIOD, and $\Delta \ln D_j^{\text{foreign}}$ is the growth of foreign final demand for sector j's output between 2001 and 2007. The identifying assumption is that pre-WTO production linkages are predetermined relative to post-WTO platform activity, and that foreign demand shocks shift the production-network linkages without directly affecting information flows except through the production network.

Table 8 reports results. The IV estimate restricts to $N = 413$ sector pairs with nonzero predicted values. The IV downstream coefficient is $\hat{\beta}_D^{IV} = 0.148$ (SE = 0.061), slightly larger than the OLS estimate of 0.113, consistent with attenuation bias in OLS rather than upward endogeneity bias. The first-stage Kleibergen–Paap F-statistic is 98.81, well above conventional thresholds. The IV upstream coefficient remains insignificant (0.040, SE = 0.061).

Table 8: Information Network Shares and Production-Network Linkages: OLS and IV

	OLS		IV	
	(1)	(2)	(3)	(4)
Down_{ij}	0.113**		0.148**	
	(0.053)		(0.061)	
Up_{ij}		−0.034		0.040
		(0.038)		(0.061)
N	413	413	413	413
Adj. R^2	0.017	−0.002	–	–
KP F-stat			98.81	239.36

Notes: Dependent variable: share of sector i's outbound knowledge directed to sector j (2007–2017). Down_{ij}: share of i's output purchased by j in 2006 WIOD. Up_{ij}: share of i's inputs sourced from j. Columns (3)–(4) instrument with pre-WTO (2001) input-output shares × post-WTO foreign demand growth. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Robust SE.

3.1 Empirical Summary and Connection to Theory

The empirical findings highlight the complex nature of information flows in international trade. While digital platforms reduce direct communication costs, firms' incentives to exchange information remain shaped by geographic proximity, production linkages, and economic complementarities. The strong relationship between information flows and supply chains suggests that information acquisition is not an independent firm decision but a strategic and network-driven process.

This insight has important theoretical implications. Existing trade models often treat information frictions as exogenous constraints. Our findings suggest that information diffusion is endogenous to supply-chain relationships: suppliers strategically share information with buyers to expand their own market opportunities. Additionally, the persistence of geographic frictions suggests that even in a digital era, physical location continues to play a crucial role. Motivated by these findings, the next section develops a model that incorporates both endogenous information acquisition and strategic knowledge sharing.

4. Model

The empirical facts point to four features that a model must accommodate. First, the platform operates through the extensive margin—more destinations, not higher prices or larger shipments per market—implying that information reduces fixed rather than variable trade costs (Fact 1). Second, adoption is unequal across regions and skewed toward large firms, requiring heterogeneous learning costs across locations and a fixed-cost adoption decision that interacts with firm productivity (Fact 2). Third, information flows decay sharply with geographic distance and concentrate along domestic supplier–buyer links, so the model needs spatially structured, network-mediated knowledge transmission (Facts 3–4). Fourth, competition shapes both the quality of information shared and the magnitude of adoption gains, motivating an oligopolistic market structure in which markups respond endogenously to information shocks.

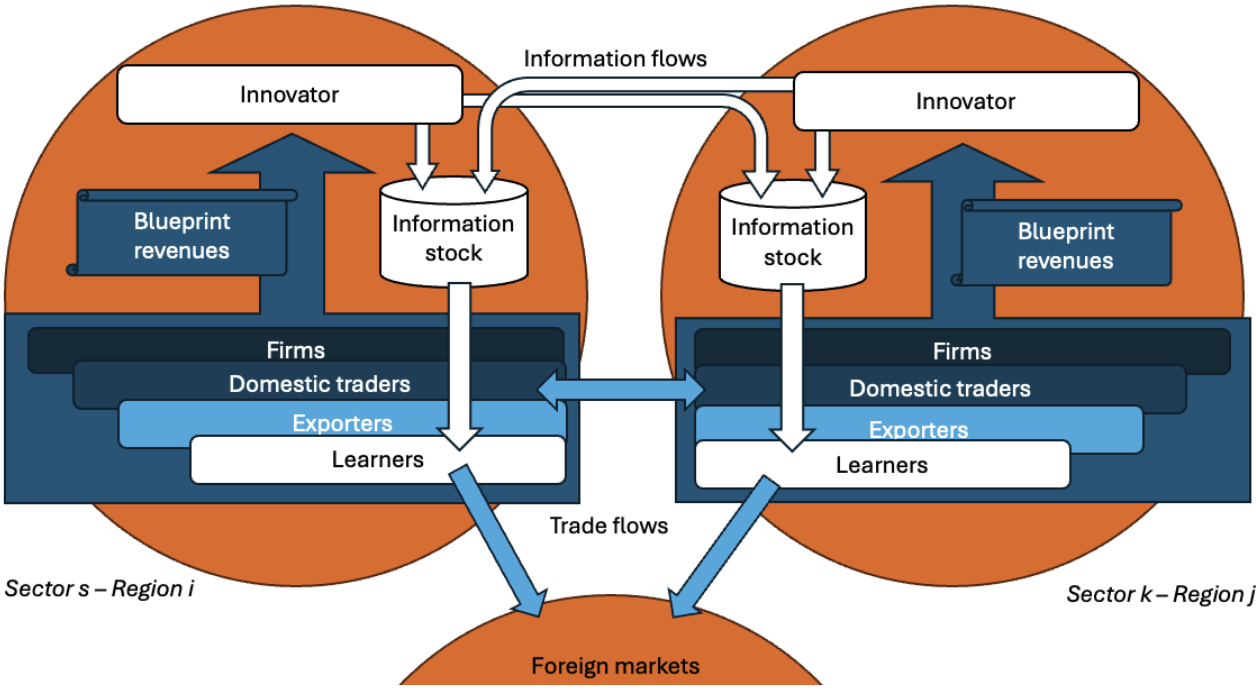
We now develop a model that incorporates all four features. The framework embeds endogenous information provision inside an oligopolistic trade model with heterogeneous

firms, variable markups, and domestic production networks. We first describe the environment, then characterize firm decisions and equilibrium, and finally derive the analytical results that connect the model to the empirical moments.

4.1 Environment and Timing

Figure 6 summarizes the model environment. The world contains N domestic locations and one foreign market x. Within each location there are S sectors. Households consume sectoral composites and supply labor inelastically. There is an exogenous measure of households L_i in each location, each supplying one unit of labor at wage w_i. Labor is immobile across locations but mobile across sectors and activities within each location.

Figure 6: Model Environment



Notes: Multiple domestic location-sectors and one foreign market. In each domestic location-sector, there is one owner and a continuum of intermediate goods producers. The owner supplies blueprints to local firms and information flows to domestic location-sectors. Information flows aggregate into information stocks. Firms endogenously choose to enter domestic/foreign markets and to learn from the information stock.

Preferences. The representative consumer in location j maximizes:

$$C_j = \prod_{s=1}^{S} C_{j,s}^{\beta_{j,s}}, \quad C_{j,s} = \left(\int_{\omega \in \Omega} c_{j,s}(\omega)^{(\sigma_s - 1)/\sigma_s} d\omega \right)^{\sigma_s/(\sigma_s - 1)} \tag{4.1}$$

where $\beta_{j,s} \geq 0$ are exogenous preference parameters satisfying $\sum_s \beta_{j,s} = 1$, $C_{j,s}$ is total consumption of sector s composite in location j, and σ_s is the within-sector elasticity of substitution. Let E_j denote total expenditure and $E_{j,s} = \beta_{j,s} E_j$ denote destination-sector expenditure.

Production. Intermediate goods are produced using both labor and intermediate inputs. The unit production cost in location-sector (i, s) is:

$$c_{i,s} = w_i^{1-\alpha_{i,s}} \prod_{k=1}^{S} P_{i,k}^{\alpha_{i,ks}} \tag{4.2}$$

where $\alpha_{i,s} \equiv \sum_k \alpha_{i,ks}$ denotes the total intermediate input share in location i, sector s, so that $1 - \alpha_{i,s}$ is the labor (value-added) share.

Entry. Each location-sector has one representative owner who supplies blueprints to local firms and posts information flows. A firm pays a fixed entry cost $c_{i,s} f_{i,s}^e$ to purchase a blueprint, draws a productivity z from the Pareto CDF $G_{i,s}(z) = 1 - (z/z_{i,s}^{min})^{-\theta}$, and produces a differentiated variety that can be used as both intermediate input and final good. Goods are traded across locations with iceberg trade costs $\tau_{ij,s} \geq 1$.

Timing. The model is static. First, owners choose blueprint mass and outbound knowledge flows. Second, knowledge flows aggregate into location-sector information stocks. Third, firms draw productivity, choose domestic entry, export participation, and platform learning. Finally, active firms compete in destination markets and households consume.

Two-level architecture. Following Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015), the model features two levels of aggregation. *Within* each destination-sector market, a finite number of active firms compete à la Bertrand, generating endogenous variable markups that depend on individual market shares (Lemma 5.1). *Across* the large number of location-sector pairs, standard CES aggregation applies because no single market is

large enough to affect aggregate price indices. The Pareto productivity distribution disciplines the mass of active firms and the gravity structure of bilateral trade shares, while the finite-firm pricing block generates the markup variation that distinguishes this model from constant-markup benchmarks. The cross-market aggregation block (GE-1 through GE-7 below) determines wages, price indices, expenditure shares, revenues, and trade balances. The within-market pricing block (Lemma 5.1) determines how individual firms' markups respond to changes in market structure induced by information shocks.

4.2 Knowledge Flows and Learning

Motivated by the prevalent social media adoption behavior in the data, we model learning as a binary technology-adoption decision following the literature (Melitz, 2003; Yeaple, 2005; Bustos, 2011; Choi and Shim, 2022). Knowledge is posted as a directed flow $H_{is,jk}$ from teacher (i,s) to learner (j,k). The total inbound stock available to learner (j,k) is:

$$H_{j,k} = \sum_{i \in \mathcal{N}} \sum_{s \in \mathcal{S}} H_{is,jk}. \tag{4.3}$$

Following Eaton and Kortum (2002) and Buera and Oberfield (2020), a firm that pays the fixed learning cost $c_{i,s} f^L_{i,s}$ draws an idea z^x from a Fréchet distribution $F_{i,s}(z^x) = \exp[-H_{i,s}(z^x)^{-\bar{\beta}}]$. The expected idea draw quality is:

$$\eta_{j,k} = \Gamma\left(1 - \frac{1}{\bar{\beta}}\right) H_{j,k}^{1/\bar{\beta}}, \qquad \bar{\beta} > 1. \tag{4.4}$$

The parameter $\bar{\beta}$ governs the efficiency at which information stock contributes to firm-level idea quality. One can also interpret $\bar{\beta}$ as the scale elasticity in the external economies of scale literature (Bartelme et al., 2025). A learner's effective iceberg trade cost on exports becomes $\tau_{ix,s}/(z\eta_{i,s})$, reducing the export price and improving competitiveness.

4.3 Finite-Firm Oligopoly and Variable Markups

The model's main departure from the standard Melitz framework is that each destination-sector market is served by a finite set of active firms. For a firm from origin i and sector s

selling to destination j, let

$$v_{ij,s}(z,T;\boldsymbol{H}) = \frac{X_{ij,s}(z,T;\boldsymbol{H})}{E_{j,s}} \tag{4.5}$$

denote the firm's expenditure share. Following Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015), the residual demand elasticity is

$$\varepsilon_{ij,s}(z,T;\boldsymbol{H}) = \left(\frac{1-v_{ij,s}}{\sigma_s} + \frac{v_{ij,s}}{\kappa_s}\right)^{-1}, \qquad 1 < \kappa_s \leq \sigma_s, \tag{4.6}$$

and the associated markup is

$$\mu_{ij,s}(z,T;\boldsymbol{H}) = \frac{\varepsilon_{ij,s}}{\varepsilon_{ij,s}-1}. \tag{4.7}$$

Small firms with negligible shares face elasticities near σ_s and charge markups near the constant-markup floor $\sigma_s/(\sigma_s - 1)$. Large firms face lower elasticities and charge higher markups. This creates a pro-competitive channel: when trade or learning compresses dominant firms' market shares, markups fall.

Define effective marginal cost as

$$m_{ij,s}(z,T;\boldsymbol{H}) = \begin{cases} \dfrac{c_{i,s}\tau_{ij,s}}{z}, & j \in \mathcal{N}, \\ \dfrac{c_{i,s}\tau_{ix,s}}{z\eta_{i,s}^{T}}, & j = x. \end{cases} \tag{4.8}$$

Prices satisfy the Bertrand pricing rule $p_{ij,s} = \mu_{ij,s} \cdot m_{ij,s}$. The key distinction from the constant-markup case is immediate: learning changes both marginal cost and markup because the firm's market share adjusts endogenously when its effective export cost falls. The export price can be written as:

$$p_{ix,s}(z,T;\boldsymbol{H}) = \mu_{ix,s}(z,T;\boldsymbol{H}) \cdot \frac{c_{i,s}\tau_{ix,s}}{z \cdot \eta_{i,s}^{T}} \tag{4.9}$$

which shows that the export price is governed by core productivity, the external economies of scale through $\eta_{i,s}$, and the endogenous markup.

4.4 Firm Policies

Each blueprint holder draws productivity z from $G_{i,s}(z)$. Destination-specific operating profits are:

$$\pi_{ij,s}(z, T; \boldsymbol{H}) = \left(1 - \frac{1}{\mu_{ij,s}}\right) X_{ij,s} - c_{i,s} f_{ij,s}, \tag{4.10}$$

and the firm's value is:

$$v_{i,s}(z; \boldsymbol{H}) = \max_{\{a_{ij,s}\}, x_{i,s}, T_{i,s}} \left\{ \sum_{j \in \mathcal{N}} a_{ij,s} \pi_{ij,s}(z, 0; \boldsymbol{H}) + x_{i,s} \pi_{ix,s}(z, T_{i,s}; \boldsymbol{H}) - x_{i,s} T_{i,s} c_{i,s} f^L_{i,s} \right\}. \tag{4.11}$$

Firms only enter a market or learn when it is profitable to do so. As shown in Proposition 5.2, this yields productivity cutoffs for domestic entry, export entry, and learning.

4.5 Unified Owner Problem

The paper's main modeling step combines the firm problem and the knowledge-provision problem into one owner problem. A representative owner in location-sector (i, s) chooses the mass of local blueprints $M_{i,s}$ and the full vector of outbound knowledge flows $\{H_{is,jk}\}$:

$$V_{i,s} = \max_{M_{i,s}, \{H_{is,jk}\}} \left\{ M_{i,s} \int v_{i,s}(z; \boldsymbol{H})\, dG_{i,s}(z) - c^H_{i,s} f^H_{i,s} \sum_{j,k} \frac{(H_{is,jk})^{\gamma^H}}{\gamma^H} - c_{i,s} f^e_{i,s} M_{i,s} \right\}, \tag{4.12}$$

where $\gamma^H > 1$ implies convex knowledge-posting costs. If blueprint creation is interior:

$$\int v_{i,s}(z; \boldsymbol{H})\, dG_{i,s}(z) = c_{i,s} f^e_{i,s}. \tag{4.13}$$

This formulation makes clear that knowledge posting is not an auxiliary activity. It is valued only because it changes the expected profits of local firms. The owner internalizes the externality that knowledge provision creates for downstream export performance. This is the microfoundation for the empirical pattern that knowledge flows align with downstream domestic supplier–buyer relationships: suppliers post more information to learner location-sectors where the resulting export expansion benefits the supplier through increased demand

for intermediate inputs.

5. Analytical Results and Equilibrium

5.1 Main Results

Lemma 5.1 (Residual elasticity and markup schedule)**.** *Under* (4.6)*, the residual demand elasticity* $\varepsilon_{ij,s}$ *is weakly decreasing in a firm's expenditure share* $v_{ij,s}$ *whenever* $1 < \kappa_s \leq \sigma_s$*. Consequently, the markup* $\mu_{ij,s}$ *is weakly increasing in expenditure share.*

Proof. See Appendix B.

Lemma 5.1 is the basic pro-competitive property inherited from Atkeson and Burstein (2008) and Edmond, Midrigan and Xu (2015). Large firms face lower elasticities and hence charge higher markups. When trade liberalization or learning compresses dominant firms' market shares, markups fall and consumers benefit from lower prices.

Proposition 5.2 (Cutoff characterization under oligopoly)**.** *There exist productivity thresholds* $\bar{z}_{ij,s}$*,* $\bar{z}_{ix,s}$*, and* $\bar{z}^L_{i,s}$ *such that: a firm serves domestic destination* j *iff* $z \geq \bar{z}_{ij,s}$*; it exports iff* $z \geq \bar{z}_{ix,s}$*; and among exporters it learns iff* $z \geq \bar{z}^L_{i,s}$*. These thresholds satisfy:*

$$\pi_{ij,s}(\bar{z}_{ij,s}, 0; \boldsymbol{H}) = 0, \tag{5.1}$$

$$\pi_{ix,s}(\bar{z}_{ix,s}, 0; \boldsymbol{H}) = 0, \tag{5.2}$$

$$\pi_{ix,s}(\bar{z}^L_{i,s}, 1; \boldsymbol{H}) - \pi_{ix,s}(\bar{z}^L_{i,s}, 0; \boldsymbol{H}) = c_{i,s} f^L_{i,s}. \tag{5.3}$$

Remark. Under constant markups, monotonicity of the learning gain in productivity is immediate because revenue scales as $z^{\sigma-1}$. Under variable markups, the learning gain also depends on how market shares and markups respond to the learner's cost reduction. In the calibrated model, the learning gain is strictly increasing in z for all parameter values we consider. We maintain this monotonicity as a regularity condition for the oligopoly extension; it is verified numerically in every equilibrium we compute.

The proof is in Appendix B. Proposition 5.2 shows that the oligopoly model preserves the selection logic of the constant-markup model but not its closed forms. Learning remains more attractive for productive firms, yet the threshold is now influenced by endogenous

market shares and markups. Under monopolistic competition, the learner share among exporters admits the closed form:

$$\bar{S}^L_{i,s} = \left(\frac{f^L_{i,s}}{f_{ix,s}} \frac{1}{\eta_{i,s}^{\sigma_s - 1} - 1} \right)^{-\theta/(\sigma_s - 1)}. \tag{5.4}$$

Equation (5.4) is the bridge between the observed regional adoption pattern and the model's location-sector learning costs: the learner share is higher when the learning cost $f^L_{i,s}$ is lower relative to the export entry cost $f_{ix,s}$, and when the expected idea draw $\eta_{i,s}$ is higher.

Proposition 5.3 (Unified owner first-order condition)**.** *Suppose the owner's choice of $H_{is,jk}$ is interior. Then:*

$$c^H_{i,s} f^H_{i,s} (H_{is,jk})^{\gamma^H - 1} = \mathcal{E}_{is,jk}(\boldsymbol{H}) \frac{\partial X_{jx,k}}{\partial \eta_{j,k}} \frac{\partial \eta_{j,k}}{\partial H_{j,k}}, \tag{5.5}$$

where $\partial \eta_{j,k} / \partial H_{j,k} = (1/\bar{\beta})\Gamma(1 - 1/\bar{\beta}) H_{j,k}^{1/\bar{\beta} - 1} > 0$, $A_{jk} \equiv \sum_{i' \in \mathcal{N}} X_{i'j,k}$ denotes total expenditure on sector-k goods in destination j (the demand index facing learner firms), and $\mathcal{E}_{is,jk}(\boldsymbol{H})$ summarizes the demand, market-share, and markup channels:

$$\mathcal{E}_{is,jk} = \underbrace{\frac{\partial V_{i,s}}{\partial A_{jk}} \frac{\partial A_{jk}}{\partial X_{jx,k}}}_{\textit{demand}} + \underbrace{\frac{\partial V_{i,s}}{\partial v_{ij,s}} \frac{\partial v_{ij,s}}{\partial A_{jk}} \frac{\partial A_{jk}}{\partial X_{jx,k}}}_{\textit{market-share}} + \underbrace{\frac{\partial V_{i,s}}{\partial \mu_{ij,s}} \frac{\partial \mu_{ij,s}}{\partial v_{ij,s}} \frac{\partial v_{ij,s}}{\partial A_{jk}} \frac{\partial A_{jk}}{\partial X_{jx,k}}}_{\textit{markup}}. \tag{5.6}$$

Remark. The three-channel decomposition provides economic intuition for how information affects upstream suppliers through the full general equilibrium. Because the mapping from learner exports $X_{jx,k}$ to teacher expenditure $E_{j,s}$ passes through the entire cross-market aggregation block, the decomposition is an interpretive framework for the numerical comparative statics rather than a closed-form analytical result. The signs of the demand and markup channels are established; the relative magnitudes are determined quantitatively.

Proposition 5.3 is the paper's central analytical object. It delivers the core mechanism: a teacher location-sector posts more information to a learner location-sector when (i) it sells more to that learner, and (ii) the learner has stronger export prospects. This is the microfoundation for the empirical Fact 4. Relative to the constant-markup benchmark, the

new term is the markup channel. Under constant markups, this third term vanishes and the knowledge-posting rule admits a closed form (see the Online Appendix, Section F).

Corollary 5.4 (Constant-markup nesting case)**.** *If $\kappa_s = \sigma_s$ for all sectors, then $\varepsilon_{ij,s} = \sigma_s$ and $\mu_{ij,s} = \sigma_s/(\sigma_s - 1)$ for every firm. The markup channel in $\mathcal{E}_{is,jk}$ disappears and the knowledge-posting rule collapses to:*

$$H_{is,jk}^{MC} = \left(\frac{\theta \Gamma(1 - 1/\bar{\beta})}{\sigma_s \bar{\beta} c_{i,s}^H f_{i,s}^H} \frac{X_{ij,s}}{E_j} X_{jx,k} Z_{j,k} \right)^{1/(\gamma^H - 1)}, \tag{5.7}$$

where $Z_{j,k} \equiv (\tilde{f}_{jx,k}/f_{j,k}^L)^{\delta_k} (\eta_{j,k}^{\sigma_k - 1} - 1)^{\delta_k} \eta_{j,k}^{\sigma_k - 2} H_{j,k}^{1/\bar{\beta} - 1} P_{x,k}^{1 - \sigma_k}$.

Corollary 5.5 (Competition suppresses same-sector knowledge posting)**.** *Within the unified-owner framework, the market-share channel of $\mathcal{E}_{is,jk}$ is strictly negative when teacher and learner operate in the same sector ($k = s$) and export to the same destination. Consequently, the owner of (i, s) posts strictly less knowledge to same-sector learners (j, s) than to otherwise identical different-sector learners $(j, k \neq s)$. The suppression is amplified under oligopoly relative to the constant-markup benchmark.*

Proof. See Appendix B.

Corollary 5.5 requires no new modeling ingredients. The competition effect is already embedded in the market-share channel of the unified owner's FOC: sharing knowledge with a same-sector learner tightens competition in the destination market, compresses the teacher's market share, and under oligopoly further compresses the teacher's markup. The owner internalizes this cost and posts less. The suppression is larger for owners with higher market share and higher markups—precisely the large, expert firms—which maps to the empirical finding in Fact 2 that expert competitors are the ones who opt out. Under constant markups the effect is present but attenuated, consistent with the broader finding that the oligopoly structure disciplines the quantitative magnitude of information-related effects.

Downstream profitability and the information buffer. The FOC (5.5) implies that optimal knowledge posting along link $(i, s) \to (j, k)$ is increasing in supplier exposure $\mathcal{E}_{is,jk}$ and in the responsiveness of learner exports to knowledge $\partial X_{jx,k}/\partial \eta_{j,k}$ (see Appendix B). Stronger downstream export profitability therefore leads to more knowledge posting, which generates

a larger welfare buffer between equilibria with and without information. The magnitude of the buffer is determined quantitatively.

5.2 Equilibrium Definition

5.2.1 Cross-Market Aggregation Block

The following system determines the aggregate equilibrium *conditional on* the within-market pricing block. Given markups $\mu_{ij,s}(v_{ij,s})$ from Lemma 5.1 and information-augmented cutoffs from Proposition 5.2, the cross-market block solves for wages, prices, trade shares, revenues, and expenditures:

A static equilibrium consists of wages $\{w_i\}$, unit costs $\{c_{i,s}\}$, price indices $\{P_{i,s}\}$, expenditures $\{E_{j,s}\}$, firm policy rules, destination revenues $\{R_{i,s}\}$, market shares, markups, blueprint masses $\{M_{i,s}\}$, and knowledge flows $\{H_{is,jk}\}$ such that: households optimize; firms solve (4.11); owners solve (4.12); information stocks satisfy (4.3); and goods and labor markets clear. The complete equilibrium system is:

(GE-1) Unit production cost: $c_{i,s} = w_i^{1-\alpha_{i,s}} \prod_{k=1}^{S} P_{i,k}^{\alpha_{i,ks}}$ for all i, s.

(GE-2) Expenditure shares:

$$\lambda_{ij,s} = \frac{((1+t_{ij,s})\tau_{ij,s}c_{i,s})^{-\theta}(R_{i,s}/c_{i,s})\chi_{ij,s}}{\sum_l((1+t_{lj,s})\tau_{lj,s}c_{l,s})^{-\theta}(R_{l,s}/c_{l,s})\chi_{lj,s}}, \tag{5.8}$$

where $\chi_{ij,s} = \xi_{ij,s}^{1-\sigma_s}$ for domestic destinations and $\chi_{ix,s} = \tilde{\xi}_{ix,s}^{1-\sigma_s}$ for the foreign market. Here $\xi_{ij,s}$ and $\tilde{\xi}_{ix,s}$ are compact constants collecting the fixed-cost terms, with the endogenous foreign market entry cost:

$$\tilde{f}_{ix,s} = \left[(\eta_{i,s}^{\sigma_s-1}-1)^{\delta_s+1}(f_{i,s}^{L})^{-\delta_s} + (f_{ix,s})^{-\delta_s}\right]^{-1/\delta_s}, \quad \delta_s \equiv \frac{\theta}{\sigma_s-1} - 1. \tag{5.9}$$

When learning costs are prohibitive ($f_{i,s}^{L} \to \infty$), the first term vanishes and $\tilde{f}_{ix,s} \to f_{ix,s}$, recovering the standard model without information.

(GE-3) Expenditure: $E_{j,s} = \beta_{j,s}(Y_j + D_j + T_j) + \sum_k \alpha_{j,sk}R_{j,k}$, where D_j is the trade deficit and $T_j = \sum_{i\neq j}\sum_s \frac{t_{ij,s}}{1+t_{ij,s}}\lambda_{ij,s}E_{j,s}$ is tariff revenue.

(GE-4) Revenue: $R_{i,s} = \sum_j \frac{\lambda_{ij,s}}{1+t_{ij,s}} E_{j,s}$, summing over all destinations j.

(GE-5) Income: $Y_i = w_i L_i = \sum_s (1 - \alpha_{i,s}) R_{i,s}$.

(GE-6) Information stocks and flows: $\eta_{j,k} = \Gamma(1 - 1/\bar{\beta}) H_{j,k}^{1/\bar{\beta}}$ with $H_{j,k} = \sum_{i,s} H_{is,jk}$, and knowledge flows given by the owner's first-order condition (5.5).

(GE-7) Welfare: Real income in location j is $W_j = E_j/P_j$, where $P_j = \prod_s P_{j,s}^{\beta_{j,s}}$. Using the Leontief inverse of the input–output matrix, the welfare expression can be written in terms of home expenditure shares, revenues, and unit costs (see the Online Appendix, Section G).

The system is solved jointly via fixed-point iteration, alternating between the goods block and the information block until convergence.

6. Calibration

The calibration preserves a tight mapping between the reduced-form evidence and the model.

Fact 1 (export response) identifies the return to information. In the firm-level export regression (Table 7, column 2), the coefficient on adoption interacted with log information stock is 0.031. In the model, the structural interpretation is $\beta_1 = (\sigma_s - 1)/\bar{\beta}$. With $\sigma_s = 5$, this implies $\bar{\beta} = 129.03$, indicating a thin-tailed idea distribution.

Fact 3 (gravity) identifies the curvature of information provision. The elasticity of knowledge posting with respect to bilateral trade intensity is $1/(\gamma^H - 1)$. Using the ratio of the distance elasticities in Table 6, Panel B ($0.338/0.672 = 0.503$), we obtain $\gamma^H = 2.988$.

Facts 3–4 discipline the knowledge network and learning-cost heterogeneity. Bilateral posting shares identify the normalized information-network weights. Location-sector learner shares identify variation in the ratio of learning costs to export opportunities. Standard trade elasticities, final-demand shares, and input–output coefficients come from the trade literature and the WIOD/China IO data.

Competition moments. The model targets a mean markup consistent with the feasible floor $\sigma_s/(\sigma_s - 1) = 1.25$ and fits concentration moments around that restriction.

Depreciation rate δ_H. In the dynamic model's steady state, the annual posting inflow equals depreciation times the cumulative stock. Computing the flow-to-stock ratio Q_t/H_t over the post-ban period 2014–2017 yields a mean ratio of 0.103 and a median of 0.105, yielding the

benchmark $\delta_H = 0.103$. We therefore label the benchmark value as data-consistent and report $\delta_H = 0.13$ (the full mature-platform estimate) as an upper robustness bound.

Learner-share adjustment speed κ_S. Under the irreversible adoption law $S_{t+1} = S_t + \kappa_S \max[S^* - S_t, 0]$, the event-study coefficients follow a geometric path $\beta_k = \beta_\infty[1-(1-\kappa_S)^{k+1}]$. Estimating (κ_S, β_∞) jointly by minimum distance against the revised event-study path ($k = 0, \ldots, 3$) of the flagship stacked specification yields $\hat{\kappa}_S = 0.36$ and $\hat{\beta}_\infty = 0.182$, with all predicted values within one standard error of the observed coefficients. Firms close approximately 36 percent of their adoption gap per year, implying 90 percent convergence in roughly five years.

Within-region learning-cost dispersion. A probit approach applied to the Q5/Q1 platform adoption gradient in 2017 (18.8% vs. 12.6%) identifies the size sensitivity parameter $\sigma_f \xi_f = 0.033$.[9] This modest gradient implies that a 10 log-unit difference in firm size shifts the adoption threshold by 0.33 standard deviations, consistent with the monotone adoption facts documented in Section 3.

The calibration produces province-sector-level outcomes for all 31 provinces $\times$ 21 sectors. Three features matter. First, learner shares are systematically higher in Coastal provinces (Guangdong, Zhejiang, Jiangsu) and lower in Western interior provinces. Second, information stocks are thickest in the most connected provinces, with the top five holding 85 percent of the national total. Third, implied learning costs rise monotonically from Coastal to West, consistent with the geographic adoption gradient documented in Fact 3.

[9]This parameter governs within-region learning-cost dispersion across firms. In the current quantitative implementation, learning costs vary across region-sectors but not across firms within a region-sector; the probit-based estimate disciplines the degree of size-dependent adoption heterogeneity documented in Fact 2.

Table 9: Calibration Strategy

Parameter	Description	Value	Identification
$\bar{\beta}$	Information-draw shape	129.03	Export response to adoption × info stock
γ^H	Curvature of knowledge posting	2.988	Ratio of distance elasticities (info/trade gravity)
σ_s	Small-firm demand elasticity	5	Trade literature
θ	Pareto shape parameter	5	Trade literature
κ_s	Large-firm elasticity parameter	4.95	Markup–share slope and concentration
$\beta_{j,s}$	Final demand shares	—	China 2012 MRIO (31 prov. × 21 sec.)
$\alpha_{i,s}$	Value-added share	0.628 (mean)	China 2012 MRIO
$\bar{S}^L_{i,s}$	Learner share targets	—	Platform and customs data
$H_{i,s}$	Information stock targets	—	Posting data
$H_{uk,is}/H_{i,s}$	Network shares	—	Bilateral posting data
Dynamic model			
δ_H	Knowledge depreciation	0.103	Flow-to-stock ratio, post-ban 2014–2017 (data-consistent; robustness: 0.13)
κ_S	Adoption adjustment speed	0.36	Minimum-distance against event-study path $k = 0, \ldots, 3$
$\sigma_f \xi_f$	Within-region cost sensitivity	0.033	Probit from Q5/Q1 adoption ratio (2017)

Notes: Information moments identified from Section 3. Competition moments follow the oligopoly baseline. Dynamic parameters δ_H, κ_S, and $\sigma_f \xi_f$ are newly calibrated from platform posting and event-study data; see calibration memo for details.

[b] Final demand shares, value-added shares, and input–output coefficients are province-sector-specific, computed from the China 2012 Multi-Regional Input–Output Table (42 sectors aggregated to 21). The mean value-added share across all province-sectors is 0.628; manufacturing sectors average 0.77–0.81 while services average 0.35–0.50.

7. Quantitative Results

The main counterfactual imposes a permanent, symmetric 40 percent ad valorem tariff on all trade between each Chinese province and the rest of the world—that is, the same 40 percent rate applies to both China's exports and its imports, uniformly across all 21 sectors. We compare two equilibria at each point in time: one in which the domestic information network operates as calibrated (firms post knowledge, information stocks accumulate, and learners benefit from the platform), and one in which information flows are suppressed—meaning we shut down all knowledge posting, set information stocks to zero, and force every firm to export without the platform's learning channel. The difference between these two equilibria is the *information buffer*: the welfare and export gains attributable to the platform's information network. Table 10 consolidates the cross-sectional findings across models, regions, and tariff levels. Table 11 traces the dynamic transition path over 12 years.

7.1 Cross-Sectional Results

Panel A of Table 10 reports the aggregate results under the oligopoly baseline. A symmetric 40 percent tariff lowers China's welfare by 32.35 percent without information and by 23.41 percent with information, yielding an information buffer of 8.94 percentage points. The export buffer is 12.35 ppt. Information flows thus offset roughly one-quarter of the welfare cost of fragmentation. The rest-of-world also benefits from China's information network, with a 3.93 ppt welfare buffer.

Panel B reports the welfare breakdown for all 31 mainland Chinese provinces. Information buffers range from 6.20 ppt (Gansu) to 13.04 ppt (Tianjin) across the 31 provinces, reflecting heterogeneity in information connectivity and learner-share baselines. The five highest-buffered provinces are Tianjin (13.04 ppt), Chongqing (12.82 ppt), Anhui (12.15 ppt), Hainan (12.14 ppt), and Fujian (11.35 ppt). Panel B illustrates this spread across selected coastal and interior provinces: Zhejiang (8.98 ppt), Jiangsu (8.93 ppt), and Henan (9.36 ppt) record mid-to-high buffers, while Sichuan (10.73 ppt) stands out among interior provinces—its large manufacturing base and dense production-network links generate an above-average information advantage. Gansu (6.20 ppt), Qinghai (7.04 ppt), and Tibet (6.92 ppt) record the smallest buffers among the selected provinces, consistent with their lower platform

connectivity and smaller learner-share baselines. Sector-level detail for all 21 stylized sectors is in the companion quantification file.

Panel C shows the tariff sensitivity. The welfare buffer declines as the tariff wedge intensifies: 15.53 ppt at 30%, 8.94 ppt at 40%, and 7.53 ppt at 50%. The export buffer peaks at 30.32 ppt (30%) and falls to 8.28 ppt (50%). The gradient is driven by the extensive margin: information lowers effective fixed export costs most when marginal exporters are near their break-even threshold, a margin that erodes as the tariff drives most exporters out of the market.

Table 10: Quantitative Results: Fragmentation Counterfactual (All 31 Provinces)

	Welfare (%)			*Exports (%)*		
	No-info	With-info	Buffer	No-info	With-info	Buffer
Panel A: China aggregate (40% tariff)						
China aggregate	**−32.35**	**−23.41**	**8.94 ppt**	**−69.74**	**−57.39**	**12.35 ppt**
RoW	−9.84	−5.91	3.93 ppt	—	—	—
Panel B: Selected provinces (40% tariff; full table in the Online Appendix)						
Selected coastal/eastern provinces						
Zhejiang	−32.45	−23.47	8.98 ppt	—	—	—
Shanghai	−31.73	−23.62	8.11 ppt	—	—	—
Henan	−32.62	−23.25	9.36 ppt	—	—	—
Guangdong	−32.01	−25.39	6.62 ppt	—	—	—
Jiangsu	−32.08	−23.16	8.93 ppt	—	—	—
Selected interior/western provinces						
Sichuan	−31.37	−20.65	10.73 ppt	—	—	—
Gansu	−32.41	−26.22	6.20 ppt	—	—	—
Qinghai	−32.22	−25.19	7.04 ppt	—	—	—
Tibet	−32.09	−25.17	6.92 ppt	—	—	—
Panel C: Tariff sensitivity (China aggregate)						
30% tariff	−28.31	−12.78	15.53 ppt	−59.29	−28.97	30.32 ppt
40% tariff	**−32.35**	**−23.41**	**8.94 ppt**	**−69.74**	**−57.39**	**12.35 ppt**
50% tariff	−35.63	−28.10	7.53 ppt	−77.52	−69.24	8.28 ppt

Notes: Variable-markup (oligopoly) model with $\kappa_s = 4.95$. "No-info" suppresses all domestic information flows; "with-info" is the active-information equilibrium. Buffer = no-info loss − with-info loss (ppt). Bold rows indicate the benchmark counterfactual. All 31 mainland Chinese provinces + RoW; labor endowments proportional to 2015 provincial GDP. Constant-markup comparison ($\kappa_s = \sigma_s$) is reported in the appendix.

The full 31-province breakdown and all 21-sector results are reported in the Online Appendix

7.2 Dynamic Transition

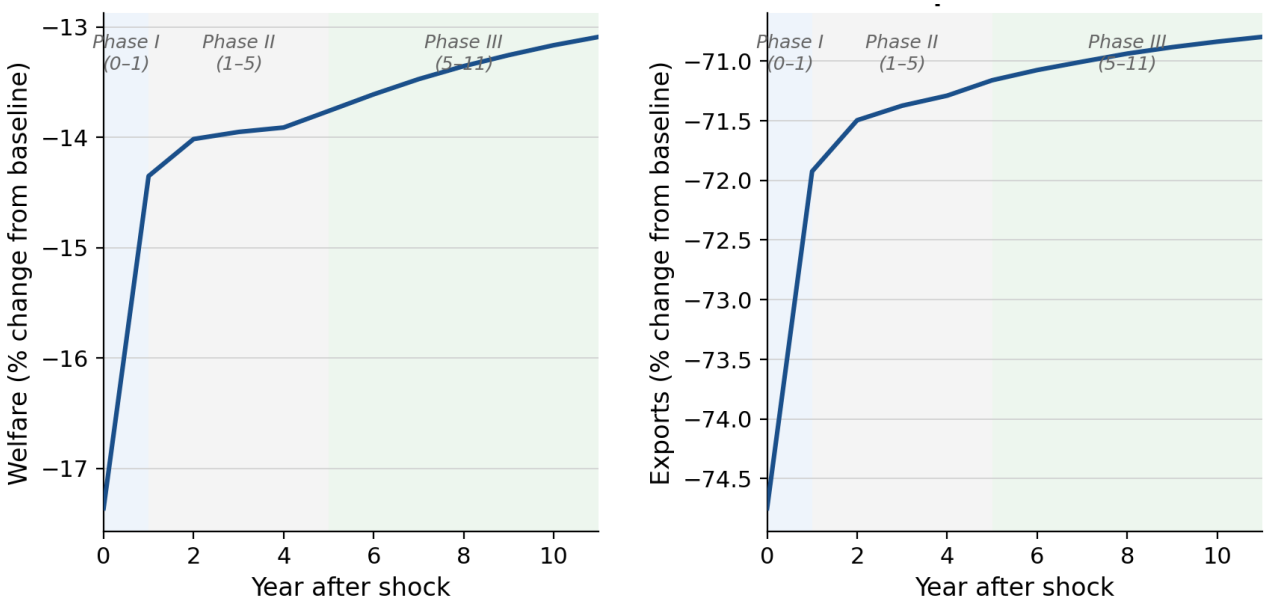
Table 11 traces the 12-year transition path. With all 31 provinces, the information stock at impact (H_0) lies substantially below its long-run steady-state target because the fragmentation shock immediately reduces the trade-driven flow of knowledge posting. From year 0 onward, the stock grows monotonically as firms adjust their posting effort toward the new equilibrium: $H_{11}/H_0 = 4.64$. The learner-share channel reinforces this: S_t/S_0 rises from 1.000 at impact to 1.078 by year 11 as firms close their adoption gap at rate $\kappa_S = 0.36$. Both channels work in the same direction—information stock accumulation and adoption expansion—so the welfare buffer amplifies continuously from 8.94 ppt at impact to 11.20 ppt at the terminal horizon, a 25 percent amplification. The export buffer similarly expands from 12.35 ppt to 15.90 ppt (+29 percent). This monotone amplification means the impact-period buffer substantially understates the long-run resilience value of the domestic information network. Learning-cost subsidies accelerate this amplification: by raising the effective adoption speed they lift the terminal buffer above 11.20 ppt without affecting the static impact. This motivates platform subsidies as a complementary resilience policy, consistent with the Hayekian insight that good institutions create decentralized mechanisms for knowledge exchange (Hayek, 1945).

Table 11: Dynamic Transition: Information Buffer Under 40% Fragmentation (31 Provinces)[10]

		State variables		*Welfare*		*Exports*	
Year	Phase	H_t/H_0	S_t/S_0	With-info	Buffer	With-info	Buffer
0 (impact)	I	1.000	1.000	−23.41%	8.94 ppt	−57.39%	12.35 ppt
1	I	3.558	1.000	−14.35%	9.94 ppt	−71.92%	14.78 ppt
2	II	4.047	1.023	−14.02%	10.28 ppt	−71.50%	15.20 ppt
5	II	4.412	1.063	−13.76%	10.53 ppt	−71.16%	15.54 ppt
7	II	4.529	1.072	−13.47%	10.82 ppt	−71.01%	15.69 ppt
9	III	4.601	1.076	−13.26%	11.03 ppt	−70.88%	15.82 ppt
11 (term.)	III	4.645	1.078	−13.09%	11.20 ppt	−70.80%	15.90 ppt

Notes: Variable-markup (oligopoly) model with $\kappa_s = 4.95$. H_t/H_0: aggregate China information stock relative to impact. S_t/S_0: aggregate learner share relative to impact. Welfare and export "with-info" are % changes from the pre-fragmentation benchmark. Buffer = no-info loss − with-info loss; the no-information path is time-invariant (−32.35% welfare, −69.74% exports). Phase I: rapid stock accumulation (years 0–1); Phase II: decelerating growth with rising learner share (years 2–7); Phase III: convergence to new steady state (years 8–11). Calibration: $\delta_H = 0.103$, $\kappa_S = 0.36$, irreversible adoption, 31 provinces.

Figure 7: Dynamic Transition Path: With-Information Welfare and Export Buffers



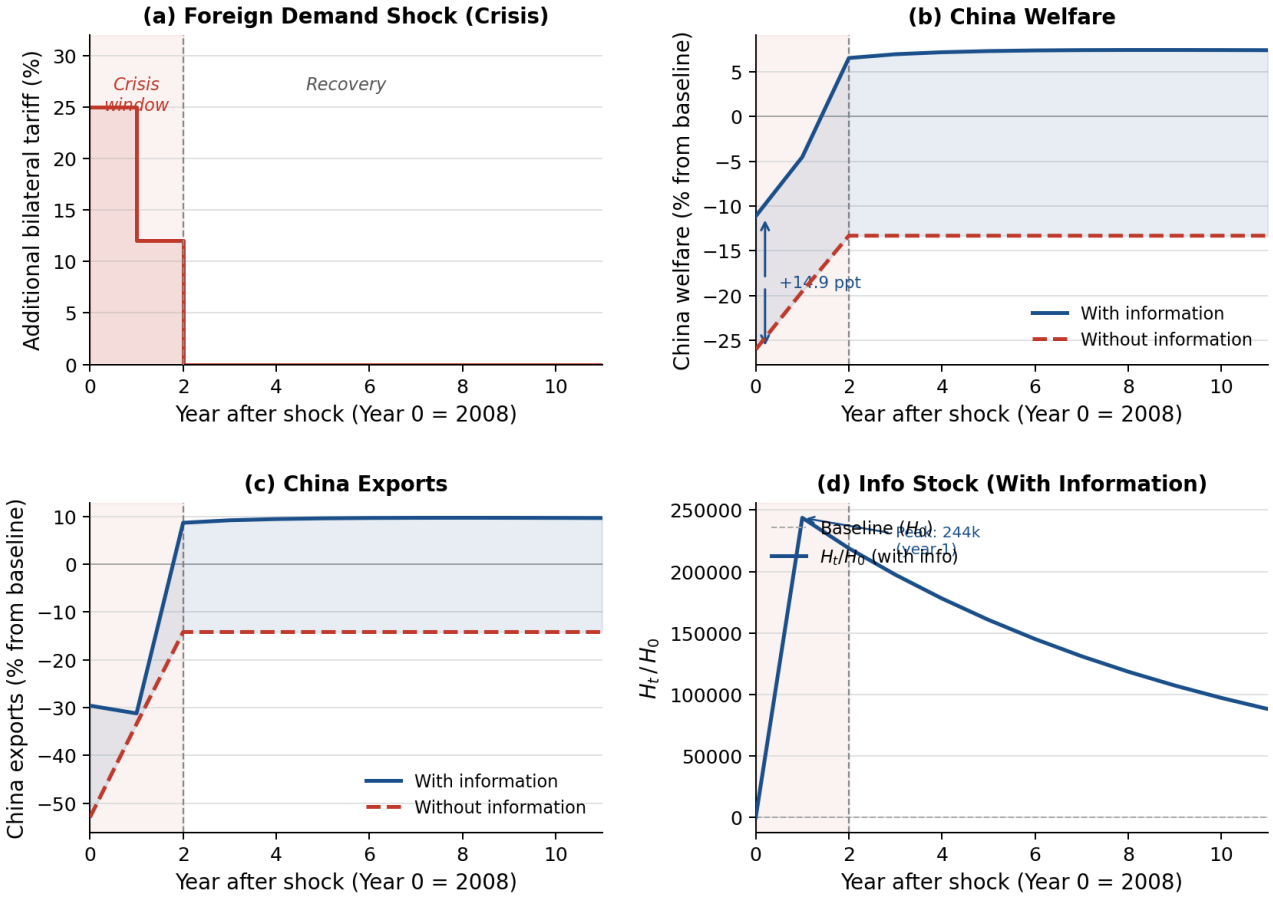
Notes: Variable-markup (oligopoly) model with $\kappa_s = 4.95$. With-information paths for welfare (left) and exports (right) relative to the pre-fragmentation baseline, under a permanent 40% bilateral tariff. Phase I (years 0–1): rapid information-stock accumulation. Phase II (years 2–7): decelerating growth as the learner-share ratchet closes the adoption gap. Phase III (years 8–11): convergence to a new steady state with welfare buffer of 11.20 ppt. Calibration: $\delta_H = 0.103$, $\kappa_S = 0.36$, all 31 mainland Chinese provinces.

7.3 Scenario: Financial Crisis (2008–09)

We model the 2008–09 global financial crisis as a temporary bilateral demand shock: a 25% tariff-equivalent in year 0, falling to 12% in year 1 and zero from year 2 onwards. This replicates the documented ~25% collapse in China's exports during 2008–09.

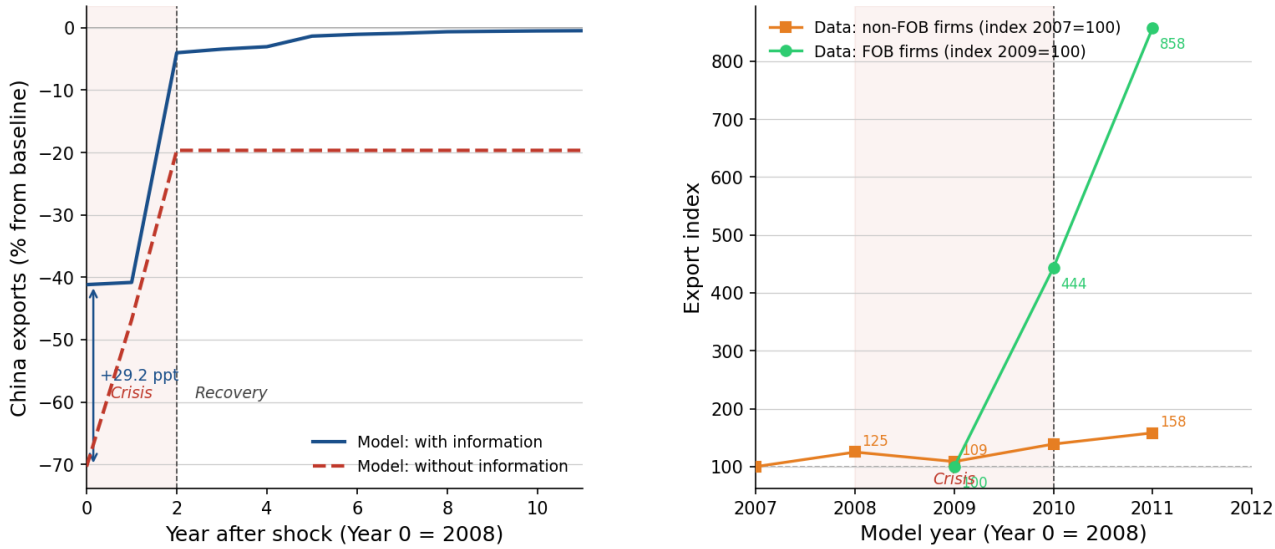
Empirical validation. Figure 9 validate the scenario against Chinese customs data. platform firms grow 343% in 2010 vs. 28% for non-platform firms, consistent with the model's predicted export buffer. New platform registrations spike 3.4-fold in 2009 (1,104 → 3,733), consistent with the model's adoption-ratchet mechanism.

Figure 8: Financial Crisis: With vs. Without Information



Notes: Variable-markup (oligopoly) model. Panel (a): tariff-shock path. Panel (b): China welfare (% from baseline); navy = with information, dashed red = without. Panel (c): China exports. Panel (d): aggregate information stock H_t/H_0. Shaded region = crisis window (years 0–2). Calibration: $\delta_H = 0.103$, $\kappa_S = 0.36$, 31 provinces.

Figure 9: Financial Crisis Validation: Export Paths vs. Customs Data



Notes: Panel (a): model export paths. Panel (b): actual customs data – non-platform exports (index 2007=100, orange) and platform exports (index 2009=100, green). Platform firms outperform non-platform by a wide margin during and after the crisis.

8. Conclusion

This paper studies how digital information networks affect trade when firms compete oligopolistically and markups are endogenous. The empirical evidence shows that platform adoption raises exports, that information remains local even online, that knowledge transmission concentrates along domestic supplier–buyer links, and that adoption is uneven across space. The model turns these facts into a quantitative system in which firms choose whether to learn and location-sector owners choose whether to post knowledge.

The quantitative results show that the information buffer survives the change in market structure. Domestic information flows reduce China's losses under fragmentation and improve export outcomes, even though gains are smaller and more uneven than in the constant-markup benchmark. The policy interpretation is direct: digital information networks are not a substitute for open trade, but they materially improve resilience when trade costs rise. Governments need not collect and distribute information directly—they can create or subsidize online communities where knowledge sharing endogenously occurs.

More broadly, the paper suggests that a complete theory of trade and information should take both network structure and market structure seriously. Digital information flows matter

for trade, but how much they matter depends on the competitive environment in which firms use them.

Supplemental Appendix

Appendix A: Empirical

A. Additional Figures Referenced from Main Text

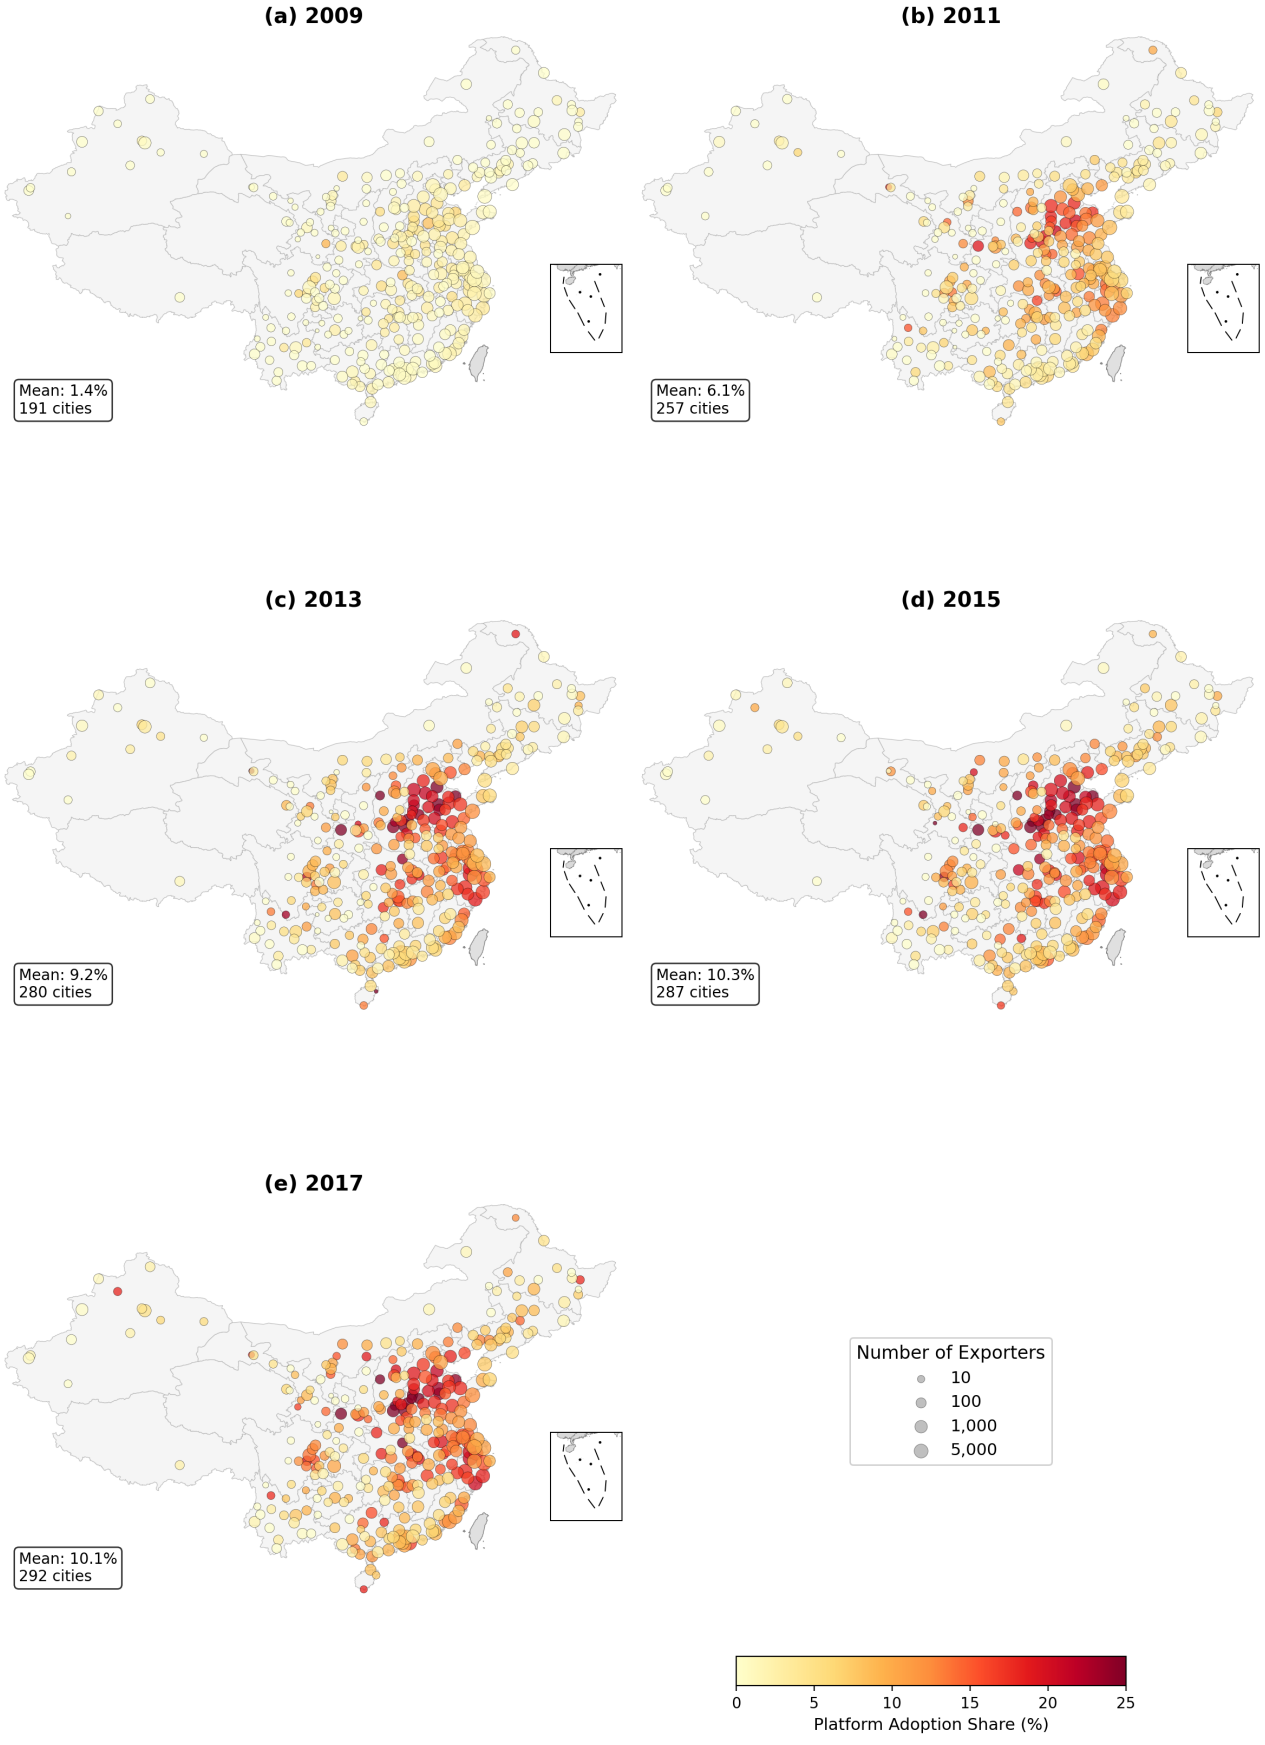
Appendix B: Theoretical

Table A1: Sample Composition Across Selection Stages

Characteristic	(1) Full customs	(2) Platform exporters	(3) Eligible pool	(4) Matched sample
Panel A: Sample size				
Number of firms	233,260	27,804	15,635	9,528
Share of full customs (%)	100.0	11.9	6.7	4.1
Panel B: Export scale (medians)				
Log export value	13.4	14.2	14.7	15.0
Destinations	3	9	12	15
HS products	3	5	7	7
Export value (USD, thousands)	660	1,425	2,496	3,369
Panel C: Export scale (means)				
Log export value	13.2	13.9	14.5	15.0
Destinations	3.9	7.8	10.1	13.3
HS products	4.1	5.9	7.4	8.5
Log unit price	3.02	3.17	3.22	3.27
Panel D: Panel structure				
Mean years observed	8.2	9.4	10.8	10.0
Median years observed	8	10	11	10
Share observed ≥ 10 years (%)	36.1	56.3	100.0	100.0
Panel E: Industry composition (%)				
Textiles & Apparel	21.3	19.5	20.0	20.7
Machinery	26.7	27.8	27.4	25.6
Misc. Manufacturing	18.0	17.2	17.0	17.8
Chemicals	11.9	13.7	13.6	13.8
Metal Products	9.5	11.0	10.9	10.8
Other	12.6	10.8	11.1	11.3
Panel F: Treatment-control balance (matched sample only)				
Treated: mean log export ($t = c-1$)	—	—	—	15.02
Control: mean log export ($t = c-1$)	—	—	—	15.04
Norm. difference (log export)	—	—	—	−0.010
Norm. difference (log dest)	—	—	—	−0.006
Norm. difference (log HS)	—	—	—	−0.005
Norm. difference (log price)	—	—	—	0.016

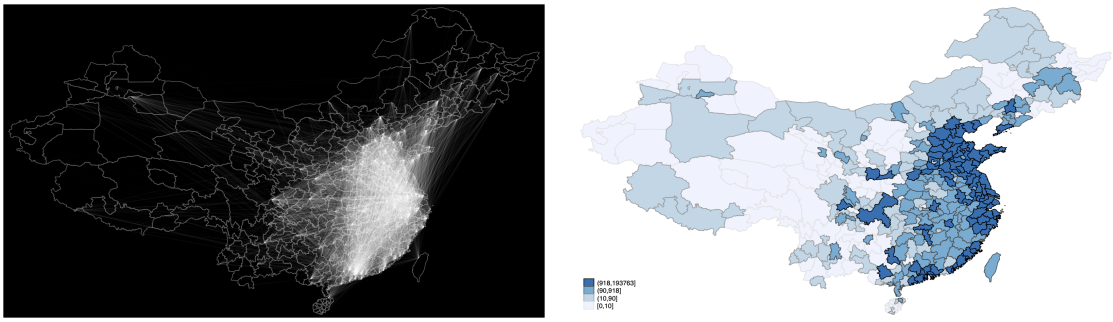
Notes: Column (1): all firms with positive exports in any year during 2000–2017 (233,260 unique firms). Column (2): customs exporters matched to the platform by firm identifier (27,804 firms). Column (3): platform exporters with ≥ 10 consecutive years of positive exports (the panel-support requirement for the $[-3, +4]$ event window plus one-year margins). Column (4): the final matched sample of 4,764 treated and 4,764 control firms after 1-NN propensity-score matching within cohort–sector cells. Panel F reports normalized differences (Imbens and Rubin, 2015) between treated and matched control firms at $t = c - 1$; values below 0.25 indicate adequate balance (Stuart and Rubin, 2008). Medians and means in Panels B–C are computed at the firm's last observed year (columns 1–3) or the pre-treatment snapshot $t = c - 1$ (column 4). Destinations and HS products are reported in levels; underlying variables in the data are in logs. Industry shares are based on the firm's customs sector classification.

Figure A1: Geographic Spread of Platform Adoption Across Chinese Cities



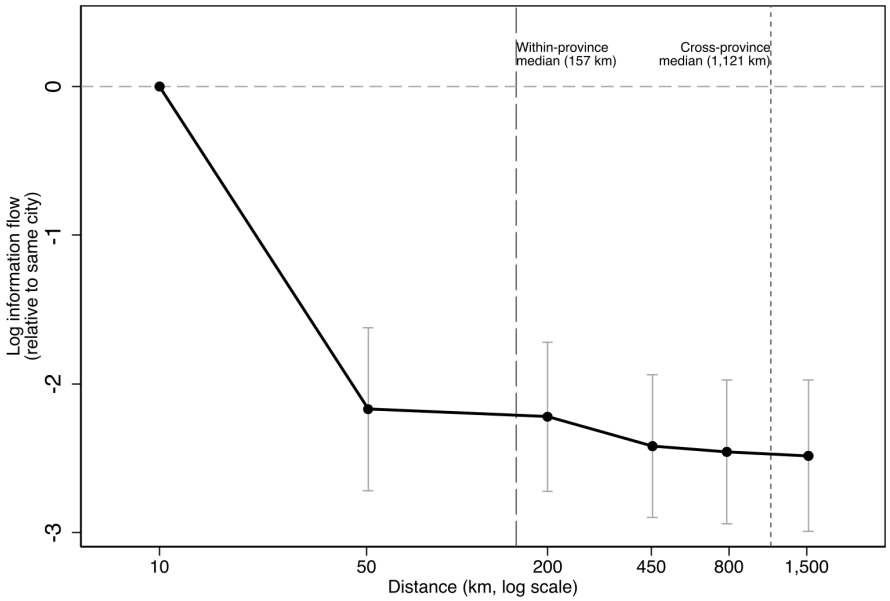
Notes: Each panel maps 339 prefecture-level cities. Color indicates adoption share (%). Bubble size proportional to log number of exporters. City-level customs data available from 2009. The boundaries shown do not imply any judgment on the legal status of any territory.

Figure A2: Geographic Distribution of Information Flows and Stocks



A. Information Flows B. Information Stocks

Figure A3: Distance Decay Profile of Information Flows



Notes: Coefficients from distance-bin dummies with source- and destination-city fixed effects.

B. Proofs

B.1 Proof of Proposition 5.2

Proof. Fix the knowledge vector $\boldsymbol{H}$ and consider destination j (domestic or export). We establish three claims.

Step 1: Operating profit is continuous and weakly increasing in z. Effective marginal cost is $m_{ij,s} = c_{i,s}\tau_{ij,s}/(z\eta_{i,s}^{T})$, which is strictly decreasing in z. Under CES demand, revenue $X_{ij,s}$ is continuous in $m_{ij,s}$ and weakly increasing in z. Expenditure share $v_{ij,s} = X_{ij,s}/E_{j,s}$ is therefore also weakly increasing in z. By Lemma 5.1, the markup $\mu_{ij,s} = \varepsilon_{ij,s}/(\varepsilon_{ij,s} - 1)$ is weakly increasing in $v_{ij,s}$ and hence weakly increasing in z. Operating profit is $\pi_{ij,s} = (1 - 1/\mu_{ij,s})X_{ij,s} - c_{i,s}f_{ij,s}$. Since both the profit margin $(1 - 1/\mu_{ij,s})$ and revenue $X_{ij,s}$ are weakly increasing in z, operating profit is weakly increasing in z and continuous by the continuity of revenue and markup in $v_{ij,s}$.

Step 2: Existence of domestic and export cutoffs. For the lowest productivity $z \to z_{i,s}^{\min}$, effective marginal cost is high, so revenue and share are small and operating profit is negative (dominated by the fixed cost $c_{i,s}f_{ij,s} > 0$). As $z \to \infty$, effective marginal cost vanishes, revenue grows without bound, and profit becomes strictly positive (finiteness of operating profits at each z is ensured by CES demand, but the limit is dominated by the variable-profit term). Since $\pi_{ij,s}(z, T; \boldsymbol{H})$ is continuous and weakly increasing in z, by the intermediate value theorem there exists a threshold $\bar{z}_{ij,s}$ solving $\pi_{ij,s}(\bar{z}_{ij,s}, 0; \boldsymbol{H}) = 0$ for each domestic destination j. We select the lowest such root and adopt a threshold-selection rule: firms enter if and only if $z \geq \bar{z}_{ij,s}$. An analogous argument yields the export cutoff $\bar{z}_{ix,s}$ from $\pi_{ix,s}(\bar{z}_{ix,s}, 0; \boldsymbol{H}) = 0$.

Step 3: Learning cutoff and its monotonicity. Consider the net gain from learning: $\Delta\pi(z) \equiv \pi_{ix,s}(z, 1; \boldsymbol{H}) - \pi_{ix,s}(z, 0; \boldsymbol{H})$. Learning sets $T = 1$, reducing effective marginal cost by the factor $\eta_{i,s}^{-1} < 1$ (since $\eta_{i,s} > 1$ when $H_{i,s} > 0$). This lower marginal cost raises the learner's revenue, expenditure share, and (by Lemma 5.1) markup relative to the non-learner with the same z. Hence:

$$\Delta\pi(z) = \left(1 - \frac{1}{\mu_{ix,s}(z, 1; \boldsymbol{H})}\right) X_{ix,s}(z, 1; \boldsymbol{H}) - \left(1 - \frac{1}{\mu_{ix,s}(z, 0; \boldsymbol{H})}\right) X_{ix,s}(z, 0; \boldsymbol{H}).$$

Both terms $(1 - 1/\mu) \cdot X$ are weakly increasing in z (Step 1), but the learner term grows weakly faster because the learner has a lower effective marginal cost and hence a weakly larger share and markup at every z. Therefore $\Delta\pi(z)$ is weakly increasing in z. Since $\Delta\pi$ is continuous and the learning cost $c_{i,s} f^L_{i,s}$ is a positive constant, there exists a threshold $\bar{z}^L_{i,s}$ solving $\Delta\pi(\bar{z}^L_{i,s}) = c_{i,s} f^L_{i,s}$. We again select the lowest root: all firms with $z \geq \bar{z}^L_{i,s}$ find learning weakly profitable. □

B.2 Proof of Proposition 5.3

Proof. Take the first-order condition of (4.12) with respect to $H_{is,jk}$. The posting-cost term contributes $c^H_{i,s} f^H_{i,s} (H_{is,jk})^{\gamma^H - 1}$ on the left-hand side. On the right-hand side, $H_{is,jk}$ enters the owner's gross value $V_{i,s}$ only through the learner's information stock $H_{j,k} = \sum_{i',s'} H_{i's',jk}$ (equation (4.3)), so $\partial H_{j,k} / \partial H_{is,jk} = 1$.

Step 1: From $H_{j,k}$ to $\eta_{j,k}$. Differentiating the idea-quality mapping (4.4):

$$\frac{\partial \eta_{j,k}}{\partial H_{j,k}} = \frac{1}{\bar{\beta}} \Gamma\left(1 - \frac{1}{\bar{\beta}}\right) H_{j,k}^{1/\bar{\beta} - 1} > 0.$$

Step 2: From $\eta_{j,k}$ to learner exports $X_{jx,k}$. Higher $\eta_{j,k}$ lowers the effective marginal cost of every learner firm in (j,k) by reducing their effective iceberg cost $\tau_{jx,k}/(z\eta_{j,k})$. Under CES demand, firm-level revenue is $X_{jx,k}(z) \propto (p_{jx,k}(z))^{1-\sigma_k} E_{x,k}$, and a lower marginal cost reduces the price $p_{jx,k} = \mu_{jx,k} \cdot c_{j,k} \tau_{jx,k}/(z\eta_{j,k})$, raising revenue for each active firm. Under oligopoly, the learner cutoff $\bar{z}^L_{j,k}$ and the endogenous markups also adjust, so the aggregate response $\partial X_{jx,k} / \partial \eta_{j,k}$ depends on the net effect across the extensive and intensive margins. We maintain $\partial X_{jx,k} / \partial \eta_{j,k} > 0$ as a maintained regularity condition, which is verified numerically in the calibrated model and holds generically when the direct cost-reduction effect dominates the equilibrium markup adjustment.[11]

Step 3: From $X_{jx,k}$ to the teacher's value $V_{i,s}$. Learner exports $X_{jx,k}$ affect the teacher location-sector (i,s) because demand for (i,s)'s goods depends on aggregate expenditure in destination j, which includes the learner sector's activity through the input–output structure.

[11] In the constant-markup limit ($\kappa_s = \sigma_s$), the sign is unambiguous because markups do not respond. Under oligopoly, the condition can in principle fail if information-induced entry compresses incumbent shares enough to reverse the sign, but this does not occur in any parameterization we consider.

Denote total expenditure on sector-k goods in destination j by $A_{jk} \equiv \sum_{i'} X_{i'j,k}$. We decompose the effect of $X_{jx,k}$ on $V_{i,s}$ through three channels:

(i) *Demand channel.* Increased learner exports raise A_{jk}, which through input–output linkages raises demand $E_{j,s}$ for the teacher's goods. The contribution is $(\partial V_{i,s}/\partial A_{jk})(\partial A_{jk}/\partial X_{jx,k})$, which is positive by the positive-linkage property established by the FOC decomposition.

(ii) *Market-share channel.* When $k = s$, learner exports in the *same* destination sector affect the teacher's expenditure share $v_{ij,s}$ through changes in the destination price index $P_{x,s}$. Additional supply lowers $P_{x,s}$, which reduces the teacher's share. The contribution is $(\partial V_{i,s}/\partial v_{ij,s})(\partial v_{ij,s}/\partial A_{jk})(\partial A_{jk}/\partial X_{jx,k})$.

(iii) *Markup channel.* Under oligopoly ($\kappa_s < \sigma_s$), the teacher's markup $\mu_{ij,s}$ depends on its expenditure share via Lemma 5.1: $\partial \mu_{ij,s}/\partial v_{ij,s} > 0$. Changes in $v_{ij,s}$ from channel (ii) therefore further affect profits through the markup adjustment. The contribution is $(\partial V_{i,s}/\partial \mu_{ij,s})(\partial \mu_{ij,s}/\partial v_{ij,s})(\partial v_{ij,s}/\partial A_{jk})(\partial A_{jk}/\partial X_{jx,k})$.

Summing these three channels yields the exposure term $\mathcal{E}_{is,jk}$ defined in (5.6). Combining Steps 1–3 by the chain rule:

$$\frac{\partial V_{i,s}}{\partial H_{is,jk}} = \mathcal{E}_{is,jk}(\boldsymbol{H}) \cdot \frac{\partial X_{jx,k}}{\partial \eta_{j,k}} \cdot \frac{\partial \eta_{j,k}}{\partial H_{j,k}}.$$

Setting $\partial V_{i,s}/\partial H_{is,jk} = c^H_{i,s} f^H_{i,s} (H_{is,jk})^{\gamma^H - 1}$ delivers equation (5.5). □

B.3 Proof of Corollary 5.4

Proof. Set $\kappa_s = \sigma_s$ in (4.6). The residual demand elasticity becomes $\varepsilon_{ij,s} = ((1 - v_{ij,s})/\sigma_s + v_{ij,s}/\sigma_s)^{-1} = \sigma_s$, independent of market share. The markup is therefore constant: $\mu_{ij,s} = \sigma_s/(\sigma_s - 1)$ for every firm. Since $\partial \mu_{ij,s}/\partial v_{ij,s} = 0$, the markup channel in $\mathcal{E}_{is,jk}$ (the third term in (5.6)) vanishes identically.

It remains to derive the closed form (5.7). Under constant markups, export revenue of a learner firm with productivity z is $X_{ix,s}(z, 1) = \tilde{\sigma}_s E_{x,s} (p_{ix,s}/P_{x,s})^{1-\sigma_s}$, where $\tilde{\sigma}_s = (\sigma_s/(\sigma_s - 1))^{1-\sigma_s}$. Integrating over the Pareto distribution of active firms, the demand

channel simplifies: $\partial V_{i,s}/\partial A_{jk}$ reduces to $\theta X_{ij,s}/(\sigma_s E_j)$ using the standard gravity form of bilateral expenditure shares. The market-share channel collapses to the revenue-weight $X_{ij,s}/E_j$ because the markup is constant and only the revenue effect remains.

Substituting into the FOC (5.5) and collecting the learner-side terms $\partial X_{jx,k}/\partial \eta_{j,k}$, which under constant markups and Pareto productivity reduce to standard gravity-trade expressions (see, e.g., Arkolakis, Costinot and Rodríguez-Clare 2012 for the Pareto integration), yields the compound $Z_{j,k}$:

$$c_{i,s}^H f_{i,s}^H (H_{is,jk}^{MC})^{\gamma^H - 1} = \frac{\theta}{\sigma_s} \frac{X_{ij,s}}{E_j} \cdot X_{jx,k} \cdot Z_{j,k} \cdot \frac{1}{\bar{\beta}} \Gamma\left(1 - \frac{1}{\bar{\beta}}\right),$$

where $Z_{j,k} \equiv (\tilde{f}_{jx,k}/f_{j,k}^L)^{\delta_k} (\eta_{j,k}^{\sigma_k - 1} - 1)^{\delta_k} \eta_{j,k}^{\sigma_k - 2} H_{j,k}^{1/\bar{\beta} - 1} P_{x,k}^{1-\sigma_k}$. Solving for $H_{is,jk}^{MC}$ delivers (5.7). □

B.4 Proof: Downstream Profitability and the Information Buffer

Proof. We establish the result in two steps.

Step 1: Optimal posting is increasing in downstream profitability. The FOC (5.5) equates the convex marginal cost $c_{i,s}^H f_{i,s}^H (H_{is,jk})^{\gamma^H - 1}$ (strictly increasing in $H_{is,jk}$ since $\gamma^H > 1$) to the marginal benefit $\mathcal{E}_{is,jk} \cdot (\partial X_{jx,k}/\partial \eta_{j,k}) \cdot (\partial \eta_{j,k}/\partial H_{j,k})$. By the positive-linkage property established by the FOC decomposition, the marginal benefit is positive for active links. By the implicit function theorem applied to the FOC, optimal $H_{is,jk}^*$ is strictly increasing in $\mathcal{E}_{is,jk}$ and in $\partial X_{jx,k}/\partial \eta_{j,k}$, since an increase in either shifts the marginal benefit upward while the marginal cost schedule is unchanged.

Step 2: Higher posting implies a weakly larger welfare buffer. The welfare buffer is defined as $\Delta W_j \equiv W_j(\boldsymbol{H}^*) - W_j(\mathbf{0})$, the difference between welfare with and without information. By the positive-linkage property established by the FOC decomposition, W_j is weakly increasing in posted knowledge along active links. Stronger downstream export profitability raises $H_{is,jk}^*$ (Step 1), which raises $H_{j,k}$ (equation (4.3)), which raises $\eta_{j,k}$ (equation (4.4)), which lowers learner export prices and raises export revenue. The resulting increase in aggregate posted knowledge weakly raises $W_j(\boldsymbol{H}^*)$ while $W_j(\mathbf{0})$ is unchanged, so ΔW_j is weakly larger. □

Proof of Lemma 5.1. Differentiate (4.6) with respect to $v_{ij,s}$: $\partial \varepsilon_{ij,s}/\partial v_{ij,s} = -((1 - v_{ij,s})/\sigma_s + v_{ij,s}/\kappa_s)^{-2}(1/\kappa_s - 1/\sigma_s)$. Because $\kappa_s \leq \sigma_s$, we have $1/\kappa_s - 1/\sigma_s \geq 0$, so $\partial \varepsilon/\partial v \leq 0$. Since

$\partial\mu/\partial\varepsilon = -1/(\varepsilon-1)^2 < 0$, by the chain rule $\partial\mu/\partial v \geq 0$. □

Proof of Corollary 5.5. Consider the market-share term in (5.6):

$$\mathcal{E}^{\text{ms}}_{is,jk} = \frac{\partial V_{i,s}}{\partial v_{ix,s}} \cdot \frac{\partial v_{ix,s}}{\partial P_{x,s}} \cdot \frac{\partial P_{x,s}}{\partial X_{jx,k}}. \tag{B.1}$$

When $k = s$, the learner (j, s) sells in the same destination sector as the teacher's firms. We sign each factor:

(a) $\partial P_{x,s}/\partial X_{jx,s} < 0$. The CES price index satisfies $P_{x,s}^{1-\sigma_s} = \sum_l P_{lx,s}^{1-\sigma_s}$. Higher learner exports raise $P_{x,s}^{1-\sigma_s}$ and hence lower $P_{x,s}$ (since $1-\sigma_s < 0$).

(b) $\partial v_{ix,s}/\partial P_{x,s} > 0$. The teacher's expenditure share satisfies $v_{ix,s} \propto P_{x,s}^{\sigma_s - 1}$ holding the teacher's own price fixed. Since $\sigma_s > 1$, the share is increasing in $P_{x,s}$.

(c) $\partial V_{i,s}/\partial v_{ix,s} > 0$. Owner value includes $(1-1/\mu_{ix,s})X_{ix,s}$, where revenue $X_{ix,s} = v_{ix,s}E_{x,s}$ is increasing in share and markup $\mu_{ix,s}$ is weakly increasing in share by Lemma 5.1.

The product is therefore $(+)(+)(-) < 0$ when $k = s$. When $k \neq s$, $\partial P_{x,s}/\partial X_{jx,k} = 0$ so $\mathcal{E}^{\text{ms}}_{is,jk} = 0$.

Under constant markups ($\kappa_s = \sigma_s$), the markup term $\partial\mu/\partial v$ vanishes, attenuating the suppression. Under oligopoly ($\kappa_s < \sigma_s$), Lemma 5.1 ensures $\partial\mu/\partial v > 0$, adding a second negative contribution that amplifies the competition effect. □

The Online Appendix, available on the authors' websites, contains: (i) complete step-by-step derivations of all theoretical results (preferences, production, pricing, cutoffs, price indices, the owner's FOC, the full GE system, and notation); (ii) the full derivation of the dynamic oligopoly model with all proofs; (iii) full quantitative results for all 31 provinces under the constant-markup benchmark.

References

Acemoglu, Daron, and Pablo D Azar. 2020. "Endogenous production networks." *Econometrica*, 88(1): 33–82.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "What do we learn from Schumpeterian growth theory?" In *Handbook of economic growth*. Vol. 2, 515–563. Elsevier.

Ahn, JaeBin, Chan Kim, Nan Li, and Andrea Manera. 2026. "Knowledge Diffusion Through FDI: Worldwide Firm-Level Evidence." *Journal of International Economics*, 160: 104205.

Akcigit, Ufuk, and William R Kerr. 2018. "Growth through heterogeneous innovations." *Journal of Political Economy*, 126(4): 1374–1443.

Allcott, Hunt, and Matthew Gentzkow. 2017. "Social media and fake news in the 2016 election." *Journal of economic perspectives*, 31(2): 211–236.

Allen, Treb. 2014. "Information frictions in trade." *Econometrica*, 82(6): 2041–2083.

Alvarez, Fernando E, Francisco J Buera, Robert E Lucas, et al. 2013. "Idea flows, economic growth, and trade." National Bureau of Economic Research.

Anderson, James E, and Eric Van Wincoop. 2004. "Trade costs." *Journal of Economic literature*, 42(3): 691–751.

Antweiler, Werner, and Murray Z Frank. 2004. "Is all that talk just noise? The information content of internet stock message boards." *The Journal of finance*, 59(3): 1259–1294.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review*, 102(1): 94–130.

Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi. 2023. "Spatial production networks." National Bureau of Economic Research.

Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-market, trade costs, and international relative prices." *American Economic Review*, 98(5): 1998–2031.

Bailey, Michael, Drew Johnston, Theresa Kuchler, Johannes Stroebel, and Arlene Wong. 2022. "Peer effects in product adoption." *American Economic Journal: Applied Economics*, 14(3): 488–526.

Bartelme, Dominick, Arnaud Costinot, Dave J Donaldson, and Andrés Rodríguez-Clare. 2025. "The textbook case for industrial policy: Theory meets data." *Journal of Political Economy*, 133(5): 1527–1573.

Bartik, Timothy J. 1991. "Who Benefits from State and Local Economic Development Policies?" W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito. 2019. "Production networks, geography, and firm performance." *Journal of Political Economy*, 127(2): 639–688.

Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. "A toolkit of policies to promote innovation." *Journal of economic perspectives*, 33(3): 163–184.

Bollen, Johan, Huina Mao, and Xiaojun Zeng. 2011. "Twitter mood predicts the stock market." *Journal of computational science*, 2(1): 1–8.

Broocks, Annette, and Johannes Van Biesebroeck. 2017. "The Impact of Export Promotion on Export Market Entry." *Journal of International Economics*, 107: 19–33.

Buera, Francisco J, and Ezra Oberfield. 2020. "The global diffusion of ideas." *Econometrica*, 88(1): 83–114.

Buera, Francisco J, and Robert E Lucas Jr. 2018. "Idea flows and economic growth." *Annual Review of Economics*, 10: 315–345.

Bui, Ha, Zhen Huo, Andrei A Levchenko, and Nitya Pandalai-Nayar. 2022. "Noisy global value chains." National Bureau of Economic Research.

Bustos, Paula. 2011. "Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms." *American economic review*, 101(1): 304–340.

Caballero, Ricardo J, and Adam B Jaffe. 1993. "How high are the giants' shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth." *NBER macroeconomics annual*, 8: 15–74.

Carballo, Jerónimo, Marisol Rodriguez Chatruc, Catalina Salas Santa, and Christian Volpe Martincus. 2022. "Online business platforms and international trade." *Journal of International Economics*, 137: 103599.

Carlino, Gerald, and William R Kerr. 2015. "Agglomeration and innovation." In *Handbook of Regional and Urban Economics*. Vol. 5, 349–404. Elsevier.

Carvalho, Vasco M. 2014. "From micro to macro via production networks." *Journal of Economic Perspectives*, 28(4): 23–48.

Carvalho, Vasco M, and Alireza Tahbaz-Salehi. 2019. "Production networks: A primer." *Annual Review of Economics*, 11(1): 635–663.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The effect of minimum wages on low-wage jobs." *The Quarterly Journal of Economics*, 134(3): 1405–1454.

Chaney, Thomas. 2008. "Distorted gravity: the intensive and extensive margins of international trade." *American Economic Review*, 98(4): 1707–21.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *Quarterly Journal of Economics*, 139(2): 891–936.

Choi, Jaedo. 2024. "From adoption to innovation: State-dependent technology policy in developing countries." *IMF Working Papers*, 2024(154): 1.

Choi, Jaedo, and Younghun Shim. 2022. "Technology adoption and late industrialization." Working Paper.

Cotterlaz, Pierre, and Arthur Guillouzouic. 2025. "The percolation of knowledge across space." *Journal of International Economics*, 153: 104026.

Cristea, Anca D. 2011. "Buyer-seller relationships in international trade: Evidence from US States' exports and business-class travel." *Journal of international economics*, 84(2): 207–220.

Disdier, Anne-Célia, and Keith Head. 2008. "The Puzzling Persistence of the Distance Effect on Bilateral Trade." *Review of Economics and Statistics*, 90(1): 37–48.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica*, 70(5): 1741–1779.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2015. "Competition, markups, and the gains from international trade." *American Economic Review*, 105(10): 3183–3221.

Ellison, Glenn, Edward L Glaeser, and William R Kerr. 2010. "What causes industry agglomeration? Evidence from coagglomeration patterns." *American Economic Review*, 100(3): 1195–1213.

Fadeev, Evgenii. 2024. "Creative Construction: Knowledge Sharing and Cooperation Between Firms." Duke University. Working paper, June 2024.

Goldman Sachs. 2023. "The creator economy could approach half-a-trillion dollars by 2027." Goldman Sachs Research.

Hajdini, Ina, Saten Kumar, Samreen Malik, Jordan J. Norris, and Mathieu Pedemonte. 2025. "Supply Chain Networks and the Macroeconomic Expectations of Firms." Federal Reserve Bank of Cleveland Working Paper 25-17.

Hausmann, Ricardo, and Dani Rodrik. 2003. "Economic development as self-discovery." *Journal of development Economics*, 72(2): 603–633.

Hayek, FA. 1945. "The Use of Knowledge in Society." *The American Economic Review*, 35(4): 519–530.

Head, Keith, and Thierry Mayer. 2014. "Gravity Equations: Workhorse, Toolkit, and Cookbook." In *Handbook of International Economics*. Vol. 4, , ed. Gita Gopinath, Elhanan Helpman and Kenneth Rogoff, 131–195. Elsevier.

Imbens, Guido W, and Donald B Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction.* Cambridge University Press.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic localization of knowledge spillovers as evidenced by patent citations." *the Quarterly journal of Economics*, 108(3): 577–598.

Jones, Charles I. 2005. "Growth and ideas." In *Handbook of economic growth*. Vol. 1, 1063–1111. Elsevier.

Keller, Wolfgang. 2002. "Geographic Localization of International Technology Diffusion." *American Economic Review*, 92(1): 120–142.

Krugman, Paul. 1991. *Geography and Trade.* Cambridge, MA:MIT Press.

Lashkaripour, Ahmad, and Volodymyr Lugovskyy. 2023. "Profits, scale economies, and the gains from trade and industrial policy." *American Economic Review*, 113(10): 2759–2808.

Liu, Ernest. 2019. "Industrial policies in production networks." *The Quarterly Journal of Economics*, 134(4): 1883–1948.

Lucas Jr, Robert E, and Benjamin Moll. 2014. "Knowledge growth and the allocation of time." *Journal of Political Economy*, 122(1): 1–51.

Marshall, Alfred. 2009. *Principles of economics: unabridged eighth edition.* Cosimo, Inc.

Melitz, Marc J. 2003. "The impact of trade on intra-industry reallocations and aggregate industry productivity." *econometrica*, 71(6): 1695–1725.

Melitz, Marc J, and Stephen J Redding. 2021. "Trade and innovation." National Bureau of Economic Research 28945.

Munch, Jakob, and Georg Schaur. 2018. "The effect of export promotion on firm-level performance." *American Economic Journal: Economic Policy*, 10(1): 357–387.

Perla, Jesse, and Christopher Tonetti. 2014. "Equilibrium imitation and growth." *Journal of Political Economy*, 122(1): 52–76.

Perla, Jesse, Christopher Tonetti, and Michael E Waugh. 2021. "Equilibrium technology diffusion, trade, and growth." *American Economic Review*, 111(1): 73–128.

Portes, Richard, and Helene Rey. 2005. "The determinants of cross-border equity flows." *Journal of international Economics*, 65(2): 269–296.

Rauch, James E. 2001. "Business and social networks in international trade." *Journal of economic literature*, 39(4): 1177–1203.

Rauch, James E, and Joel Watson. 2002. "Entrepreneurship in international trade." National Bureau of Economic Research 8708.

Rhee, Yung W. 1989. *Small trading companies and a successful export response: lessons from Hong Kong.* World Bank Industry and Energy Department, Policy, Planning and Research.

Rosenblat, Tanya S, and Markus M Mobius. 2004. "Getting closer or drifting apart?" *The Quarterly Journal of Economics*, 119(3): 971–1009.

Sampson, Thomas. 2016. "Dynamic selection: an idea flows theory of entry, trade, and growth." *The Quarterly Journal of Economics*, 131(1): 315–380.

Startz, Meredith. 2021. "The value of face-to-face: Search and contracting problems in Nigerian trade." Working Paper.

Stuart, Elizabeth A, and Donald B Rubin. 2008. "Best Practices in Quasi-Experimental Designs: Matching Methods for Causal Inference." In *Best Practices in Quantitative Methods*. , ed. Jason W Osborne, 155–176. Sage Publications.

Theodosiou, Marios, and Evangelia Katsikea. 2013. "The export information system: An empirical investigation of its antecedents and performance outcomes." *Journal of International Marketing*, 21(3): 72–94.

Tian, Lin, and Yue Yu. 2023. "Geographic Integration and Firm Exports Evidence from China." Working Paper.

Volpe Martincus, Christian, and Jerónimo Carballo. 2008. "Is Export Promotion Effective in Developing Countries? Firm-Level Evidence on the Intensive and the Extensive Margins of Exports." *Journal of World Business*, 43(2): 111–122.

Volpe Martincus, Christian, and Jerónimo Carballo. 2012. "Export promotion activities in developing countries: What kind of trade do they promote?" *The Journal of International Trade & Economic Development*, 21(4): 539–578.

Wei, Shang-Jin, Ziru Wei, and Jianhuan Xu. 2021. "On the market failure of "missing pioneers"." *Journal of Development Economics*, 152: 102705.

Yeaple, Stephen Ross. 2005. "A simple model of firm heterogeneity, international trade, and wages." *Journal of international Economics*, 65(1): 1–20.