Learning by AI: Market Intelligence and Exporting*

Alexis Antoniades
Georgetown University

Chuan He
CSUST

Zheming Liang
PKU

Mingzhi (Jimmy) Xu
PKU

February 3, 2026

Abstract

Early adoption of artificial intelligence (AI) reshaped how firms responded to market dynamics by enhancing data collection and analysis. Linking China's universe of customs shipments to millions of online job ads, we tracked AI hiring in sales, marketing, and analytics to build a firm-level proxy for non-production AI and map its exposure across products within firms. To structure our analysis, we introduce a model in which firms confront asymmetric information about heterogeneous consumer preferences and show that AI mitigates these frictions by sharpening firms' ability to learn demand patterns across markets. The model predicts—and the data confirm—that AI-intensive firms fine-tune their product mix and prices with greater precision: they are more likely to export, expand their product lines, and adjust market choices. Crucially, this refinement appears as narrower price dispersion but wider quantity dispersion across destinations—effects strongest for differentiated goods, larger firms, and sales to high-income economies. Together, the evidence shows that AI eases demand-side information frictions, allowing firms to optimize their global reach strategies.

Keywords: artificial intelligence; export behavior; product differentiation

JEL Classification: F14, O14, J24

*We thank Tania Babina, Rosario Crinò, Emmanouil Chatzikonstantinou, Lorenzo Caliendo, Muhammed Yildirim, Kristina McElheran, Hong Ma, Ran Jing, Yuliang Li, Ningyuan Jia, and seminar and conference participants at IEFS Annual Conference (China), the Italian Trade Study Group Conference, the CESifo Workshop on the Economic and Societal Impacts of AI, and the NCER–CCER Seminar on the Chinese Economy for helpful comments. Mingzhi Xu acknowledges financial support from the National Natural Science Foundation of China (Grant Nos. 72322007 and 72573001). All authors contributed equally and are listed alphabetically. The authors declare no competing financial interests. All views expressed are those of the authors and do not necessarily reflect those of their institutions, their executive directors, or the countries they serve. Corresponding author: Chuan He (chhe@csust.edu.cn), Changsha University of Science and Technology. All remaining errors are our own.

1 Introduction

We use data on Chinese exporters to examine how early adoption of artificial intelligence (AI) reshaped exporters' behavior. Understanding these effects is important given the transformative impact that AI can have on business strategy. Industry reports and firm case studies suggest that AI-assisted forecasting outperforms traditional methods: prediction errors fall by roughly 20–50 percent, sales growth rises by about 10–15 percent, and inventory costs decline by 20–30 percent.[1] In China, the diffusion of AI-enabled business analytics in foreign trade is also reflected in industry accounts. One prominent example is Alibaba's international e-commerce platform, which uses AI to provide sellers with market intelligence on consumer trends across more than 200 overseas markets (China Trade Post, 2023).[2]

By making information processing, communication, and prediction faster, cheaper, and more accurate, AI has the potential to relax a central constraint in international trade: costly information frictions.[3] These frictions limit firms' ability to infer demand conditions, identify profitable market niches, and tailor products and marketing to heterogeneous foreign buyers. In this sense, AI may strengthen learning from exporting by improving how firms collect and interpret market-relevant information, thereby reducing the costs of operating in unfamiliar markets and widening the performance gap between adopters and non-adopters.[4]

To explore how AI adoption affects exporters, we merge China's universe of customs declarations with millions of online job vacancy postings (OJV) from 2014–2016. This window is dictated by customs-data availability and, like Aghion et al. (2025) for France, captures an early stage of diffusion, well before large language models became widespread. Importantly, it is also a period in which firms were already deploying AI in economically meaningful ways, so

[1] Singla et al. (2025) and Raskovich et al. (2025) provide cross-country evidence on how AI is changing business strategy and value creation. Ozturk (2024) and Pansare (2024) compile sectoral evidence on AI-powered forecasting and performance, including retail supply chains, food & beverage, fast fashion, and global logistics.

[2] Figure A.1 reproduces our translation of the report's description of these demand-intelligence and marketing functions.

[3] Brynjolfsson (1997) emphasizes how information technology can improve decision-making by lowering the costs of processing and communication. Jensen (2007) shows that price dispersion can reflect information frictions and that improved communication technologies (mobile phones) can raise market efficiency. Agrawal et al. (2022) and Marwala et al. (2023) stress AI's comparative advantage in generating more accurate and timely predictions, in part by extracting signals from large, noisy, and incomplete datasets.

[4] By *information frictions*, we mean demand-side barriers, including heterogeneous preferences and imperfect knowledge of market conditions, that limit firms' effectiveness abroad (see Manova and Zhang, 2012; Allen, 2014; Eaton et al., 2021). Our view is complementary to Allen (2014): rather than modeling learning as a fully dynamic sequence, we treat AI as an added technology that reduces informational barriers within a period. The interpretation is also related in spirit to Arkolakis (2008), where firms incur increasing marginal marketing costs to reach additional customers; here, firms can invest in AI to better understand and serve both existing and potential foreign customers.

our data cover an active early-adoption phase rather than a pre-adoption baseline.[5] Focusing on forecasting and market-analytics AI is also advantageous because its effects can translate into exporting decisions over relatively short horizons through information-processing margins such as targeting, pricing, and inventory planning.[6]

Following Babina et al. (2024), we proxy AI adoption with AI-oriented hiring. We classify a vacancy as AI-related if its text contains at least one term from a curated dictionary (e.g., *artificial intelligence*, *machine learning*, *deep learning*, and synonymous Chinese phrases), and aggregate these postings to construct a firm-year measure of AI capability. We then translate this firm-year measure into a *firm–product–year* exposure measure by allocating a firm's AI capability across its products using each product's export sales share measured in the pre-adoption period. The premise is that market-analytics AI is data-intensive: products with richer transaction histories generate more training data, so AI is more likely to affect decisions for products with larger prior sales. Out of 3,149,469 online job vacancies posted by 178,358 firms (including both exporters and non-exporters) during this period, we find that 4.6% of all postings contain at least one AI-related keyword. Moreover, 25.5% of firms posted at least one AI-related vacancy. Among the subset of exporters—firms that can be matched to customs data—we find that 6.3% of their postings reference an AI-related skill, and 30.6% of these firms advertise at least one AI-related hire.

Unlike automation technologies and industrial robots, which are primarily adopted to improve production and reduce costs, firms typically use AI for predictive analytics, data mining, planning optimization, and pricing optimization.[7] It is precisely because of these non-production applications that we expect AI to alleviate information frictions in export markets—providing the central motivation for our analysis. To explore this mechanism, we

[5]Evidence discussed by Babina (2025) points to a sharp rise in AI-related hiring beginning around 2014 (e.g., postings referencing "ML" terms), while later waves of AI adoption become prominent only after 2020 (e.g., "agent" -style applications). This timing supports our focus on prediction and market-analytics uses of AI in the mid-2010s.

[6]In contrast, industrial AI in production settings often requires complementary capital and organizational reconfiguration, so measured productivity effects may be delayed; see Mcelheran et al. (2025), which documents "productivity J-curves" for industrial AI in manufacturing.

[7]A large and growing literature examines how automation affects firms and markets by altering production technologies. For example, Acemoglu and Restrepo (2018, 2019) study the effects of automation on productivity and employment, while Acemoglu and Restrepo (2022) link it to wage inequality. Agrawal et al. (2023) examine AI investment across industries to assess its role in economic growth via production, and Babina et al. (2024) argue that AI fosters growth through product innovation (building on frameworks such as Hottman et al. (2016) and Argente et al. (2020)). Other research investigates the impact of robotics on global value chains and task-based international specialization (Faber, 2020; Stapleton and Webb, 2020; De Backer and DeStefano, 2021; Krenz et al., 2021; Artuc et al., 2023). Our contribution is to shift attention from production technologies to *non-production* tasks. We show that AI in market-facing activities (e.g., demand forecasting and marketing) reshapes firms' pricing and quantity allocation across destinations, even without changing marginal costs. Figure A.2 echoes this view in managers' accounts.

identify a subset of AI-related job postings that reference tasks consistent with these applications. For example, a job posting requiring knowledge of *machine learning* and *tensorflow* is flagged as an AI-related vacancy. If, in addition, the job description includes terms *market*, *sell*, *customer*, *business*, *predictive analytics*, or *predictive model*, we also classify the posting as an AI-Sales vacancy.[8] Among all firms in our sample, we find that 9.68% posted at least one AI-Sales job. Among exporters, this share rises slightly to 10.85%.

Next, we examine how AI-Sales hiring—which we use as a proxy for AI adoption in terms of sales—affects exporters' behavior. To structure our analysis, we first introduce a model of heterogeneous firms facing asymmetric information about consumer demand. In this framework, firms do not fully observe consumer preferences for their products when deciding production quantity. Acquiring such information requires a costly technology—namely, AI—that allows firms to uncover the true parameter of consumer tastes and mitigate information frictions. Firms in the model therefore make two key decisions: the quantity to supply, which determines the market clearing price, and whether to adopt AI. We begin by analyzing the model in a closed-economy setting and then extend it to incorporate international trade across multiple countries.

The model yields three core predictions. First, larger and more competitive firms are more likely to adopt AI. Second, AI adoption leads to *lower price dispersion* and *higher quantity dispersion* for a given product across export destinations. Third, AI adoption would result in more market adjustments. The intuition behind these predictions stems from that AI reduces information frictions. In the absence of reliable demand information, firms base their quantity supply decisions on average expected preferences. This leads them to under-supply *strong* markets—those with higher-than-average demand—and over-supply *weak* markets. By adopting AI, firms gain more accurate knowledge of market-specific demand, allowing them to reallocate quantities more efficiently: increasing supply to strong markets and reducing supply to weak ones. As a result, the variance in quantities across markets rises. Prices, by contrast, move in the opposite direction: they fall in strong markets (due to greater supply) and rise in weak markets (due to reduced supply), leading to a decline in price dispersion. This reallocation improves allocative efficiency, which is reflected in a lower aggregate price index. The gains from AI adoption are therefore larger in environments with greater demand uncertainty and heterogeneity.

[8] In the robustness section, we consider an alternative way to identify AI–Sales hiring based on SOC-coded occupations. Under this approach, a vacancy is classified as sales-related if it falls into 11-2011 (Advertising and Promotions Managers), 11-2021 (Marketing Managers), 13-1161 (Market Research Analysts and Marketing Specialists), or 41-xxxx (Sales and Related Occupations, including supervisors, sales representatives, and retail sales workers). We also show that the results are not sensitive to whether *predictive analytics* and *predictive model* are included in the sales keyword list.

We take these predictions to the data using our firm–product–year AI exposure measure, which weights a firm's AI capability toward products with richer pre-adoption transaction histories. This mapping yields a triple-difference design: we compare the post-adoption change for data-rich versus data-poor products within adopters and subtract the same contrast within non-adopters. Identification therefore comes from whether adoption induces a larger differential shift for products with more training data, after netting out firm-wide shocks and common product- or destination-level time effects. Consistent with the theory, exporters that post AI-related vacancies are larger and export a broader set of HS 8-digit products to more destinations than non-adopters.[9] Among adopting exporters, those posting AI-Sales vacancies are larger still and operate on broader product and destination margins. Our main finding speaks directly to the model's mechanism: AI adoption increases the cross-destination dispersion of quantities for a given product while reducing the dispersion of prices. This joint pattern is consistent with AI lowering demand-side information frictions and improving the allocation of sales across export markets.

To further strengthen identification against endogenous AI adoption, we complement our fixed-effects design with an IV strategy and robustness checks. The IV uses a policy-driven supply shifter for AI talent: a firm–year measure of local access to AI-trained graduates, constructed from nearby universities' *lagged* admissions quotas in AI-related majors and weighted by inverse distance. Firms closer to universities with larger quotas face lower recruiting frictions and are more likely to hire AI–Sales workers, while quotas reflect centralized, long-run workforce planning rather than firm-specific short-run conditions. The results are also stable across a wide range of checks, including excluding major municipalities and trimming outliers, addressing uneven vacancy coverage, varying the AI–Sales definition (including SOC-based classification) and keyword lists, controlling for non-sales AI, broader data-technology hiring, and traditional sales hiring, redefining dispersion measures to net out product-year and product-destination components, and running placebo tests with randomized adoption. Across these exercises, the central pattern remains stable.

We also explore heterogeneity across firms, products, and destinations. The effects are more pronounced for larger exporters, for differentiated products, and for exports to developed markets. This pattern is consistent with the model's logic. Larger firms generate richer transaction histories and can more readily deploy data-intensive tools, so AI has greater scope to improve demand inference and market allocation. Information frictions are also inherently more severe for differentiated goods, where preferences are harder to pin down and successful exporting relies more on market-specific targeting. Finally, developed markets

[9] Aghion et al. (2025) reports similar evidence for France, finding that AI adopters are larger and more productive.

tend to feature finer demand segmentation and stronger returns to accurate forecasting and adaptation, making AI-enabled learning particularly valuable in those destinations.[10]

Related Literature This paper builds on a growing literature examining the impact of artificial intelligence on firm behavior. While recent macro-level studies fail to find strong evidence of AI on productivity so far (Brynjolfsson and Unger, 2023), micro-level evidence suggests otherwise. For example, Babina et al. (2024) show that AI adoption—measured through AI hiring intensity—is associated with labor reallocation, productivity gains, and higher markups, driven largely by product innovation. Similarly, Alekseeva et al. (2021) find that firms hiring for AI roles grow faster and are more likely to innovate. Using French firm-level survey data, Aghion et al. (2025) report that AI-adopting firms are larger, more productive, and more skill-intensive, with higher post-adoption productivity. Czarnitzki et al. (2023) reach similar conclusions using German data, finding that AI improves firms' information-processing capabilities and productivity. Our paper contributes to this literature by offering new empirical evidence on the effects of AI adoption specifically among exporters. In contrast to existing studies that focus primarily on production-side outcomes, we develop a theoretical framework and empirical strategy to show how AI enhances *allocative efficiency* in non-production tasks—especially through better demand forecasting and sales optimization.[11]

More broadly, this study builds on research using online job vacancy (OJV) data to assess the impact of AI adoption across domains beyond firm performance. Scholars have used such data to examine labor market dynamics (Deming and Kahn, 2018; Felten et al., 2019; Acemoglu and Restrepo, 2022; Antoniades et al., 2025b), the financial sector (D'Acunto et al., 2019; Grennan and Michaely, 2020; Abis and Veldkamp, 2024), entrepreneurship (Gofman and Jin, 2024), the demand for STEM education (Deming and Noray, 2019; Antoniades et al., 2025c), and even political outcomes (Magistro et al., 2025; Antoniades et al., 2025a; Balcazar, 2023). A key distinction of our work is its focus on the adoption of AI for *sales-related* and *non-production* activities, which is less explored in the existing literature.

Finally, we contribute to a rich literature on information frictions and learning-by-exporting as key determinants of trade performance. Exporting requires information about foreign demand, and information asymmetries can inhibit firm participation and success in international markets (Manova and Zhang, 2012; Allen, 2014; Raj and Seamans, 2018; Gold-

[10] See Beraja et al. (2023) on scale and data advantages in AI adoption, Åkerman et al. (2022) on more severe information frictions for differentiated goods, and Van Biesebroeck (2005); De Loecker (2007) on stronger learning-by-exporting effects in developed markets.

[11] Babina et al. (2025) study how AI adoption reduces second-moment volatility in firm sales, earnings, and cash flows for U.S. firms. However, they do not offer a mechanism to explain these effects. Our work helps fill this gap by explicitly modeling the role of AI in mitigating information frictions.

farb et al., 2019b; Eaton et al., 2021; Ferencz et al., 2022; Macedoni, 2022). Allen (2014) attributes up to half of regional price dispersion to information frictions. Hottman et al. (2016) and Munch and Nguyen (2014) highlight how firm-market demand heterogeneity explains much of sales dispersion. More recently, Kim and Kim (2024) show that information frictions account for 40% lower export volumes and 15% higher export concentration in Korean data. Bai et al. (2020), using data from AliExpress, demonstrate how search frictions significantly affect seller performance in global e-commerce. Improving information flow thus enhances trade outcomes: Steinwender (2018), using the historical rollout of the telegraph in 1866, finds that improved communication reduced price levels and volatility while increasing trade volumes. Relatedly, Åkerman et al. (2022) show how broadband adoption in Norway affected trade patterns, while Lodefalk et al. (2022) study the role of export credit guarantees in Sweden in mitigating cross-border frictions.

In the learning-by-exporting literature, several studies show that export experience reduces uncertainty and leads to performance gains. Bastos et al. (2018) find that price volatility declines with export experience. Cross-country evidence also supports this view: Van Biesebroeck and Zaurino (2025) show stronger learning effects when exporting to advanced economies in Sub-Saharan Africa; De Loecker (2007) report similar findings using data from Slovenia; and Iacovone and Javorcik (2010) show that firms learn to drop low-performing products and shift toward higher-quality offerings. Our contribution lies in highlighting AI as a new channel that facilitates learning-by-exporting and reduces information frictions. We document new stylized facts on how AI adoption affects the dispersion of prices and quantities across export destinations.

The remainder of the paper is structured as follows. Section 2 describes the data. Section 3 presents the theoretical framework. Section 4 outlines the empirical methodology, Section 5 the baseline results, and Section 6 the robustness checks. We conclude in Section 7.

2 Data and Variable Description

2.1 Data Sources

We study the impact of AI hiring on exporters' behavior by combining Chinese customs data with data on online job vacancies (OJV) posted by these firms during the same period. The OJV data are scraped from the website of a leading recruitment and human resource services platform in China, beginning in 2014. This job portal offers a comprehensive collection of postings, disaggregated by city, industry, and job type. It covers all four centrally administered municipalities, 293 out of 333 prefecture-level administrative regions, and 59

industries across China.

Each posting reports structured information on wage offers (as a range), non-wage benefits, and required education and experience, and also includes unstructured text describing job tasks, firm background, and additional requirements (e.g., language proficiency and computer skills). Appendix C illustrates the format of these postings (Figure A.6). For a detailed description of the dataset, see He et al. (2021). Related studies that use Chinese online vacancy data include Kuhn and Shen (2013) and Fang et al. (2020).

We merge the OJV data to firm-level customs transactions from the General Administration of Customs of China using firm names common to both datasets. The customs data report exports by product (HS 8-digit), destination, value, and quantity. Because the customs data end in 2016, our sample covers 2014–2016. To ensure meaningful cross-destination variation in firm–product prices and quantities, we drop firms with zero reported export quantities (anomalous) and firms serving only one destination. The resulting sample includes 177,015 exporters in 2014 and 327,044 in 2016. We successfully match 4.1% and 3.9% of exporters to the OJV data in 2014 and 2016, respectively. Despite the modest match rate, matched firms account for 12.8% and 21.2% of all vacancies and for 13.9% and 15.2% of total export value in 2014 and 2016. Appendix A assesses representativeness in detail.

2.2 The AI Inputs Measure

Following a growing literature that uses online job vacancy (OJV) data to study the impact of AI on labor markets and firm-level dynamics (Goldfarb et al., 2019a; Alekseeva et al., 2021; Acemoglu et al., 2022; Babina et al., 2024), we measure AI-related hiring by identifying the presence of specific keywords associated with AI in job postings. The list of keywords—such as *machine learning*, *artificial intelligence*, and *supervised learning*—is provided in Table A.1.[12] A job posting is classified as AI-related if it contains at least one keyword from this list. We define a dummy for firm-level AI adoption as 1 if it posts at least one AI-related vacancies up to the current year, and 0 otherwise.

Given our focus on the role of AI in non-production activities—such as marketing, pricing strategy, and customer analytics—we further identify a subset of AI-related postings that reference non-production tasks. We label this measure *AI-Sales*, defined as the dummy

[12]Several keywords in our list capture widely used components of the modern AI stack, even when they do not include "AI" in their name. For example, Hadoop provides the data infrastructure for storing and processing large-scale datasets that often underpin machine-learning deployment; Figure A.7 illustrates this link. Likewise, Apache Spark is a standard distributed computing framework for large-scale data processing and includes a widely used machine-learning library (Spark ML) for model training. Some technical terms do not admit a one-to-one translation into Chinese, which can add noise to a keyword-based classifier. We therefore construct translations based on standard industry usage and re-estimate our main specifications after excluding the potentially ambiguous translated terms. The results are unchanged.

of a firm's AI-related job postings that also include keywords *sales*, *marketing*, *customer*, *business analytics*, *predictive analytics*, or *predictive models* in the job description.[13] Table 1 provides descriptive statistics for the full sample. Between 2014 and 2016, our dataset contains 3,149,469 job postings by 178,358 firms. On average, 4.6% of postings mention at least one AI-related skill, while 1.4% of postings mention both AI and a non-production-related term. At the firm level, 25.5% of firms post at least one AI-related vacancy, and 7% advertise AI-related jobs for non-production roles.

We further disaggregate the data by firm size and industry.[14] Large firms post the highest share of AI-related jobs—8% of all their vacancies—and more than half (56%) of them post at least one AI-related position. Across industries, the highest AI Intensity is found in Scientific Research and Technology Services (7.6%), Information Transmission, Software, and IT Services (6.7%), and a residual "Other Industry" category (5.4%). In contrast, Transportation, Storage, and Postal Services (0.4%), Construction and Real Estate (2.4%), and Leasing, Business Services, and Finance (2.8%) exhibit the lowest intensity. Panel B of Table 1 focuses on the subset of 31,167 firms matched to the customs data. These exporters show a higher AI Intensity (6.3%) compared to the full sample. Notably, even within the short time frame of our sample, the growth in AI-related hiring is more pronounced among exporters—increasing from 5.7% in 2014 to 6.7% in 2016—than among all firms (from 4.4% to 4.7% over the same period).

The data reveal that, as expected, exporters tend to be larger and more likely to utilize AI than non-exporters. To better understand heterogeneity within exporters, we examine whether there are systematic differences between those that employ AI talent and those that do not. Table 2 provides descriptive statistics for four groups: exporters without AI hires (Column 1), exporters with AI hires (Column 2), exporters with AI hires excluding AI-Sales (Column 3), and exporters that specifically hire for AI-Sales roles (Column 4).

We find that, among exporters, firms hiring AI talent are significantly larger, export a broader range of products, and serve more foreign markets. Furthermore, within the subset of AI-adopting exporters, those that hire for AI-Sales positions tend to be larger and more diversified in terms of products and destinations than those hiring AI talent more generally. However, these differences are not statistically significant.

[13] As a robustness check, we employ an alternative approach that uses occupation-level classifications. Specifically, we flag any AI-related job post that belongs to a non-production occupation (e.g., marketing, advertising, sales) as AI-Sales. We discuss this approach and present corresponding results in the robustness section.

[14] The full list of industries is presented in Table A.2.

Table 1: Data Description: AI-Related Job Vacancies and Hiring Firms in China, 2014–2016

	Number of firms N (1)	Online job vacancies OJVs (2)	AI-intensity AI (3)	AI-intensity sales AI-sales (4)	Share of firms hiring AI AI-firms (5)	Share of firms hiring AI-Sales AI-sales-firms (6)
Panel A: All firms (Online Job Vacancy Data)						
Overall						
All Years	178,358	3,149,469	4.6%	1.4%	25.5%	7.0%
2014	56,870	832,875	4.4%	1.3%	25.2%	6.7%
2015	61,370	1,108,852	4.5%	1.4%	25.4%	7.0%
2016	60,118	1,207,742	4.7%	1.4%	25.9%	7.2%
By Company scale (employees), all years						
<50	98,069	1,111,279	4.1%	1.4%	19.2%	5.3%
50-150	36,310	605,736	4.5%	1.4%	27.4%	7.5%
150-500	28,969	665,926	5.1%	1.5%	34.4%	9.1%
500-1,000	8,176	303,282	4.6%	1.1%	41.4%	11.0%
1,000-5,000	5,958	364,706	4.7%	1.2%	48.7%	13.8%
5,000-10,000	464	53,588	5.1%	0.9%	47.6%	17.9%
>10,000	412	44,952	8.0%	2.6%	56.3%	22.6%
By industry, all years						
Transportation, storage and postal services	5,916	85,059	2.4%	0.4%	16.0%	2.7%
Accommodation and catering; Wholesale and retail	43,320	599,556	3.3%	1.5%	17.9%	6.0%
Information transmission, software and information technology services	13,827	337,816	6.7%	1.6%	38.8%	11.8%
Public administration; Health; International organization; Electricity, heat, gas and water	1,338	17,198	4.9%	1.1%	27.3%	5.1%
Agriculture, forestry, animal husbandry and fishery; Mining	875	8,893	2.8%	0.7%	15.7%	2.5%
Manufacturing	50,167	794,756	5.3%	1.5%	27.8%	6.5%
Residential services, repair and other services; Education; Culture, sports and entertainment	4,107	62,727	3.2%	0.9%	20.1%	4.4%
Scientific research and technology services	22,050	413,170	7.6%	2.3%	37.6%	11.6%
Leasing and business services; Finance	20,925	495,160	2.8%	0.9%	21.4%	5.7%
Construction; Real estate	14,688	315,218	2.4%	0.6%	21.0%	4.9%
Other industry	1,145	19,916	5.4%	0.9%	27.0%	5.7%
Panel B: Subset of firms from customs data matched with OJV data						
Overall						
All Years	31,167	578,714	6.3%	1.9%	30.6%	7.8%
2014	7,254	115,817	5.7%	1.5%	29.7%	6.9%
2015	11,148	203,584	6.1%	1.8%	29.7%	7.6%
2016	12,765	259,313	6.7%	2.0%	31.9%	8.4%
By Company scale (employees), all years						
<50	10,897	119,470	4.5%	1.3%	18.4%	4.5%
50-150	6,787	91,438	5.0%	1.5%	26.4%	6.8%
150-500	7,895	145,951	6.2%	2.0%	35.8%	8.6%
500-1,000	2,959	81,605	6.2%	1.5%	45.2%	10.7%
1,000-5,000	2,311	108,488	8.3%	2.5%	59.1%	16.7%
5,000-10,000	167	10,949	7.0%	1.4%	61.7%	23.4%
>10,000	151	20,813	11.6%	4.4%	70.2%	31.1%
By industry, all years						
Transportation, storage and postal services	387	7,440	2.6%	0.3%	18.6%	4.1%
Accommodation and catering; Wholesale and retail	5,379	82,453	4.9%	2.1%	21.2%	6.6%
Information transmission, software and information technology services	547	28,914	9.6%	3.1%	57.6%	23.2%
Public administration; Health; International organization; Electricity, heat, gas and water	43	1,195	13.1%	6.3%	44.2%	14.0%
Agriculture, forestry, animal husbandry and fishery; Mining	120	1,690	4.3%	1.5%	21.7%	4.2%
Manufacturing	21,378	367,333	5.9%	1.6%	30.4%	6.6%
Residential services, repair and other services; Education; Culture, sports and entertainment	118	4,952	7.1%	1.1%	42.4%	11.0%
Scientific research and technology services	2,385	61,169	8.9%	2.9%	46.8%	16.5%
Leasing and business services; Finance	479	15,997	5.5%	2.2%	39.7%	12.5%
Construction; Real estate	278	5,463	8.9%	2.2%	30.9%	7.2%
Other industry	53	2,108	3.0%	0.5%	35.8%	9.4%

Note: This table reports the summary for AI use in job postings and among firms. AI-intensity means the share of AI job postings among all posts, while AI-intensity sales means the share of AI-sales job postings among all posts. As some industries are small in sample, we aggregate some of them here (and only here, i.e., we still use the disaggregated industry classification reported in Table A.2 in fixed effects regarding industry).

Table 2: Data Description: Exporter Characteristics by AI Adoption Status and Sales-AI Hiring

	Chinese Exporters			
	AI vs Non-AI		Non Sales vs Sales AI	
	(1) non-AI	(2) AI	(3) non-Sales AI	(4) Sales AI
Number of Firms	21,633	9,534	7,117	2,417
Firm Size:				
p25	17	68	67	73
p50	80	232	230	240
p75	244	643	616	747
Products	33	40	40	42
Markets	16	20	20	22
Avg # Products Per Market	3	3	3	3
Avg Sales Per Product	1,715,785	2,019,746	2,202,278	1,482,269

Notes: This table reports the summary for exporters. Avg # Products Per Market is the average number of products at each market. Avg Sales Per Product is the average sales for each product.

3 Theoretical Framework

To frame the analyses that follow and make the operative channels explicit, we develop a heterogeneous firm model with asymmetric information about consumer demand to better understand the implications of AI adoption for exporting firms. Our focus is on the role of AI in enhancing business operations—such as predictive analytics and pricing optimization—rather than on production-side benefits (as would be the case with automation technologies).

A central feature of the model is the assumption that consumer preferences are private information—known only to consumers—which creates a key informational friction. Firms must decide how to allocate output across markets despite this uncertainty. We introduce AI as a costly technology that allows firms to learn about these hidden preferences, thereby mitigating information asymmetries. By incorporating AI into their decision-making, firms can better tailor their offerings across markets. The returns to AI adoption increase with the heterogeneity of market demand—that is, the richer the distribution of consumer preferences, the more valuable AI becomes. This framework enables us to analyze how AI adoption improves firms' ability to allocate output more efficiently, enhances market targeting, and ultimately affects competition in settings with imperfect information.

3.1 Closed Economy

We begin by describing a closed economy. A continuum of consumers, normalized to measure one, derives utility from consuming differentiated product varieties. The economy is populated by a large number of monopolistically competitive firms, indexed by ω. Each firm

produces a unique product variety, differentiated by quality, where firm-specific quality is denoted by $s(\omega)$ and acts as a demand shifter. In addition to firm quality, consumer demand is shaped by heterogeneous tastes for each product. We denote the consumer-specific preference for variety ω as $z(\omega)$, which also acts as a demand shifter. Thus, variation in both $s(\omega)$ and $z(\omega)$ determines firm-level demand. Labor is the only factor of production and wage is normalized to one. Free entry into the market is allowed, and in equilibrium, expected profits are driven to zero by competition.

Consumers Consumers derive utility from consuming differentiated products according to a CES utility function, with a quality shifter as in Kugler and Verhoogen (2012); Hallak and Sivadasan (2013); Feenstra and Romalis (2014); Macedoni (2022). Specifically, the utility function is given by:

$$U = \left[\int_{\omega\in\Omega} \Big(z(\omega)s(\omega)q(\omega)\Big)^{\frac{\sigma-1}{\sigma}} d\omega\right]^{\frac{\sigma}{\sigma-1}}$$

where $q(\omega)$ represents the quantity of variety ω consumed, and $s(\omega)$ captures the quality of variety ω. We assume $s(\omega)$ is a fixed, firm-specific attribute that does not vary across consumers or markets, representing vertical differentiation.

A key distinction of our model relative to prior literature is the introduction of an additional demand shifter, $z(\omega)$, which reflects consumers' horizontal tastes. This term varies across consumers and destinations, allowing for greater flexibility in modeling heterogeneity. Finally, $\sigma > 1$ denotes the elasticity of substitution, and Ω is the set of producers in the market. The demand function for variety ω is given by:

$$q(\omega) = \frac{z(\omega)^{\sigma-1}s(\omega)^{\sigma-1}p(\omega)^{-\sigma}}{P^{1-\sigma}}$$

where P is the price index, defined as:

$$P \equiv \left[\int_{\omega\in\Omega} \left(\frac{p(\omega)}{z(\omega)s(\omega)}\right)^{1-\sigma} d\omega\right]^{\frac{1}{1-\sigma}}$$

so that $\frac{p(\omega)}{z(\omega)s(\omega)}$ represents the demand-shifter-adjusted price. It follows that indirect utility is given by $U = P^{-1}$.

Firms There is a continuum of firms, each utilizing labor as the input for production. Products differ along two dimensions: quality, denoted s, and style, represented by the horizontally differentiated variable $z(\omega)$, which varies across consumers and captures taste heterogeneity. Firms incur a fixed entry cost f_e to draw a quality level s from a continuous

distribution $F(s)$, with density $f(s)$ supported on $(0, +\infty)$. Upon observing their draw, firms decide whether to enter. If they do, they pay a fixed production cost f and choose a production level q. Marginal cost is increasing in quality and given by $c = s^\alpha$, with $0 < \alpha < 1$. This reflects that higher-quality production is more labor-intensive but still provides a competitive advantage. Consumers' preference for product ω is determined by a taste shifter z, drawn independently from a continuous distribution $G(z)$ with support on $[z_l, z_h]$. This leads to heterogeneous demand responses across consumers, even for the same product.

Asymmetric Information on Demand In this framework, we assume that information about consumer preferences is private and known only to consumers. While firms offer vertically differentiated varieties—where higher-priced goods reflect higher quality—consumers exhibit heterogeneous tastes along a horizontal dimension. However, firms are unaware of these idiosyncratic tastes before entering the market.

To capture this asymmetric information, we model firms' decision-making as a sequential process. First, firms choose their production quantities based on the expected distribution of consumer preferences, without knowing the specific demand realizations they will face. Upon entering the market, firms observe the true value of consumer tastes and subsequently set prices to clear the market and maximize ex-post profits.

This sequential structure—production first, pricing second—not only enhances analytical tractability but also reflects real-world operational constraints. In many industries, firms must commit to production and distribution plans under substantial demand uncertainty, with limited flexibility to adjust afterward. For instance, automobile manufacturers often set production targets months in advance, relying on forecasts, and later use discounts or marketing campaigns to adapt to realized demand. This timing assumption is also standard in the trade literature, including Arkolakis et al. (2018) and Sager and Timoshenko (2024), as a means of modeling firm behavior under informational frictions.

To ensure tractability and the existence of a well-defined equilibrium, we impose mild regularity conditions on the distributions of product quality and consumer taste. Specifically, we assume that the relevant moments of both s and z exist:

Assumption 1 *The random variables s and z have finite $(\sigma - 1)$-th moments, i.e.,*

$$\mathbb{E}_s[s^{\sigma-1}] < \infty \quad \textit{and} \quad \mathbb{E}_z[z^{\sigma-1}] < \infty.$$

These conditions ensure that aggregate demand and pricing decisions remain well-defined in equilibrium and allow us to derive closed-form solutions for firm behavior.

Following the procedure outlined earlier, the model can be solved using backward induction. After the firm observes the realized value of consumer demand, $z(\omega)$, it chooses the optimal price to maximize revenue, taking the previously chosen quantity as given. The firm thus solves a standard revenue maximization problem, conditional on both the quantity decision and the observed demand. The detailed derivation of the optimal pricing rule is provided in Appendix B.1.1.

$$p(s,z,q) = (zsP)^{\frac{\sigma-1}{\sigma}} q^{-\frac{1}{\sigma}}$$

The price is set so that the goods market clears. The optimal price depends on the product's quality s, the consumer's idiosyncratic taste z, and the quantity q produced by the firm before observing demand. With the price setting rule, we move to the earlier stage where firms make their production decisions—prior to learning the realized demand taste—to maximize expected profits. These expected profits equal expected revenues net of variable production costs and fixed operating costs:

$$\max_q \ \mathbb{E}_z\Big(p(s,z,q) - c\Big)q - f$$

By substituting the ex-post optimal price into this objective, we solve for the ex-ante optimal quantity. The full derivation is provided in the appendix. We then can also calculate the optimal price ex post. The closed-form expressions are:

$$p(s,z) = \frac{\sigma}{\sigma-1} \cdot \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)} s^{\alpha}, \qquad q(s) = \left[\frac{\sigma-1}{\sigma} \cdot s^{\frac{\sigma-1}{\sigma}-\alpha} \cdot P^{\frac{\sigma-1}{\sigma}} \cdot \mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right]^{\sigma} \tag{1}$$

The expectations exist by Assumption 1, since $\sigma - 1 > \frac{\sigma-1}{\sigma}$. Notably, the firm's output decision depends on the expected value of consumer tastes, $\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)$, rather than exact tastes, highlighting the implications of asymmetric information. In the absence of precise knowledge of individual preferences, firms must base production decisions on average expected demand. This induces a misallocation: firms may overproduce goods with lower demand and underproduce those with stronger market appeal. As a result, asymmetric information leads to inefficient product-level allocations in equilibrium—too many low-demand goods, too few high-demand goods.

The ex-post profit can be expressed as:

$$\pi(s,z) = \left(\frac{\sigma}{\sigma-1}\frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)} - 1\right)\left[\frac{\sigma-1}{\sigma}\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} s^{(\sigma-1)(1-\alpha)} - f.$$

Next, to compute the expected profit the firm faces when it decides whether to stay in the

market —before knowing the demand and making the quantity decision—we integrate this ex-post profit function over all possible consumer tastes. The expected profit is given by:

$$\mathbb{E}_z(\pi(s,z)) = \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f, \quad s_p = \frac{f^{\frac{1}{(\sigma-1)(1-\alpha)}}}{((\sigma-1)P)^{\frac{1}{1-\alpha}}} \left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{-\frac{\sigma}{(\sigma-1)(1-\alpha)}} \tag{2}$$

where s_p represents the quality threshold at which the marginal firm earns zero expected profits. We show the derivation in the appendix. The equilibrium cutoff is determined by the condition of zero expected profits, as implied by free entry, $\int_{s_p}^{+\infty} \Big(\mathbb{E}_z(\pi(s,z))\Big) dF(s) = f_e$. This gives the following equation for the equilibrium cutoff implicitly:

$$\int_{s_p}^{+\infty} \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f dF(s) = f_e.$$

With the equilibrium value of the quality cutoff, we can derive the remaining key equilibrium objects—the mass of entrants, the aggregate price index, and the representative consumer's ex-post utility. These derivations are in Appendix B.1.4.

The equilibrium in the model with asymmetric information about consumer preferences is characterized by several conditions. Consumers maximize utility subject to a budget constraint. Firms maximize profits but earn zero expected profits due to free entry. In addition, both the goods and labor markets must clear. The full set of equilibrium conditions is detailed in Appendix B.1.4, where we also show that the equilibrium can be characterized by solving for a set of variables: the firm-level output decision $q(s)$, the price schedule $p(s,z)$, the cutoff quality level s_p, the mass of entrants N_e, and the aggregate price index P. In Appendix B.1.5, we formally prove the existence and uniqueness of equilibrium under asymmetric information about consumer demand.

3.2 Artificial Intelligence Technology

We next examine the role of AI technology in business operations. As highlighted by both background analysis and anecdotal evidence, the core function of AI deployed by firms in our sample revolves around leveraging advanced techniques and accumulated consumer data to better understand customer behavior, refine product offerings, and optimize production plans.[15] By applying machine learning algorithms to these data, firms can uncover emerging trends and evolving consumer preferences. This data-driven approach enables more accu-

[15] A key strength of AI in sales lies in its ability to process vast amounts of consumer data—ranging from purchasing behavior and search patterns to demographic information (Bajari et al., 2015).

rate demand forecasts, allowing firms to better align production and inventory management with market needs. To simplify the analysis—and without loss of generality—we model AI technology as a tool that enables firms to observe consumer preferences in advance, thereby allowing them to make production decisions with full information about demand. This eliminates the information frictions that would otherwise arise from demand uncertainty.

Adopting AI, however, entails substantial upfront costs, including investment in digital infrastructure, the development of tailored algorithms, and the integration of AI into existing workflows. Once implemented and internal capabilities are built, the marginal cost of applying AI to additional production decisions is relatively low. Reflecting this cost structure, we treat AI adoption—including expenditures on AI-related talent—as a fixed cost, largely independent of short-run output. Let f_{AI} denote this fixed cost. To ensure that our theoretical framework matches empirical patterns—namely, that AI adoption remains concentrated among larger firms—we introduce the following assumption.

Assumption 2 *The fixed cost of AI adoption satisfies:*

$$\frac{f_{AI}}{\mathbb{E}_z(z^{\sigma-1}) - \left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}} > \frac{f}{\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}}.$$

Note that $\mathbb{E}_z(z^{\sigma-1})$ exists by Assumption 1. The condition ensures that AI adoption is costly enough to generate selection: not all firms adopt the technology, and only a subset—typically those with higher quality or larger scale—find it optimal to do so.

Firm Behavior with AI Firms equipped with AI have full information about consumer preferences for their products, denoted by z, prior to making production decisions. This informational advantage allows them to optimize output based on precise demand signals. Consequently, their profit-maximization problem yields the following optimal production quantity and price:

$$q^{AI}(s,z|a) = \left[\frac{\sigma-1}{\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha}(zP)^{\frac{\sigma-1}{\sigma}}\right]^{\sigma}, \quad p^{AI}(s,z|a) = \frac{\sigma}{\sigma-1}s^{\alpha}, \quad \forall z.$$

With complete information, the firm treats consumer preferences as exogenously known. The optimal price becomes a constant markup over marginal cost, as in the canonical Melitz framework, and does not vary with z. The benefit of AI lies instead in output flexibility: firms adjust production dynamically in response to heterogeneous demand. When a product is more favored by consumers, firms produce more; when demand weakens, production is scaled back.

In equilibrium, this results in uniform pricing across consumer types and production quantities that vary systematically with consumer tastes:

$$p^{AI}(s,z|a) = \frac{\sigma}{\sigma-1}s^{\alpha}, \quad q^{AI}(s,z|a) = \left[\frac{\sigma-1}{\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} z^{\sigma-1}. \tag{3}$$

Compared to the benchmark without AI, adoption results in more stable prices and significantly lower price dispersion across different realizations of z. Non-AI firms must set output based on expected demand, which leads to mismatches in supply and demand and necessitates price adjustments. In contrast, AI-enabled firms can precisely tailor output to demand conditions and thus rely less on price flexibility to clear the market. This tighter alignment results in price uniformity across destinations and increased dispersion in quantities allocated across markets, reflecting more efficient reallocation. AI thereby enhances firms' ability to implement targeted export strategies and mitigates the inefficiencies stemming from asymmetric information. In effect, AI acts as a demand-revealing technology that improves allocative efficiency and competitiveness in global markets.

The profit net of fixed costs for a firm adopting AI is given by

$$\pi^{AI}(s,z \mid a) = \frac{1}{\sigma^{\sigma}}\left[(\sigma-1)s^{1-\alpha}zP\right]^{\sigma-1} - f - f_{AI}.$$

Aggregating over all possible realizations of consumer demand tastes z, we obtain the expected profit for firms using AI:

$$\mathbb{E}_z\left(\pi^{AI}(s,z \mid a)\right) = \frac{\mathbb{E}_z(z^{\sigma-1})}{\sigma^{\sigma}}\left[(\sigma-1)s^{1-\alpha}P\right]^{\sigma-1} - f - f_{AI}.$$

Since adopting AI entails an additional fixed cost, not all firms find it profitable to use AI as a tool to mitigate demand uncertainty. As shown in Appendix B.2.1, only firms with sufficiently high product quality choose to integrate AI into their production processes. The quality threshold for AI adoption, denoted by s_a^{AI}, is determined by the following condition:

$$s_a^{AI} = s_p^{AI}\left(\frac{f_{AI}\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}}{f\left(\mathbb{E}_z(z^{\sigma-1}) - \left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}\right)}\right)^{\frac{1}{(\sigma-1)(1-\alpha)}} > s_p^{AI}. \tag{4}$$

Under Assumption 2, the threshold for adopting AI exceeds the threshold for simple market entry, i.e., $s_a^{AI} > s_p^{AI}$, as formally proven in Appendix B.2.2. This result reflects the empirical observation that AI adoption is concentrated among higher-quality (and typically larger) firms. Given this, the expected profit for firms that adopt AI (i.e., those with $s \in$

$(s_a^{AI}, \infty))$ is:

$$\mathbb{E}_z(\pi^{AI}(s, z \mid a)) = \left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f + \left[\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f_{AI}.$$

For firms that do not adopt AI (i.e., $s \in (s_p^{AI}, s_a^{AI})$), their profits follow the non-AI specification discussed earlier.

In the presence of AI, the equilibrium quality cutoff s_p^{AI} is determined by the zero expected profit condition under free entry. That is,

$$\int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f \, dF(s) + \int_{s_a^{AI}}^{+\infty} \left[\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f_{AI} \, dF(s) = f_e.$$

Given the relationship between s_p^{AI} and s_a^{AI}, we can jointly solve these equations to determine the unique equilibrium thresholds. The equilibrium in the scenario where AI technology is available builds upon the case without AI but introduces additional complexity. Specifically, AI generates a new threshold along the quality spectrum: only firms with sufficiently high product quality can afford to adopt AI, while others continue to operate without it. We defer the formal definition of equilibrium to Appendix B.2.3, and establish the existence and uniqueness of equilibrium in Appendix B.2.4.

Importantly, the availability of AI leads to a reallocation of activity toward higher-quality firms. In Proposition 1, we show that the introduction of AI raises the quality threshold for firm survival relative to the benchmark case without AI, where firms make production decisions under incomplete information.

Proposition 1 *With the introduction of AI technology, the quality threshold for firm survival increases compared to a scenario in which firms cannot learn consumer preferences and must treat demand as ex ante uncertain. Formally, the equilibrium quality cutoff satisfies* $s_p^{AI} > s_p$.

The proof of Proposition 1 is presented in Appendix B.2.5. This result highlights the *selection effect* introduced by AI adoption. By enabling firms to anticipate demand more precisely, AI enhances the returns to scale from quality and sharpens competitive pressures. Firms with higher-quality products are better positioned to benefit from AI-driven demand insights, while lower-quality firms are more likely to exit. This selection mechanism raises the market participation threshold, reallocates resources toward higher-quality firms, and leads to improved aggregate efficiency in the economy.

3.3 AI in an Open Economy

We extend the model to an open-economy environment to study how demand uncertainty interacts with exporting, and how AI changes firms' export decisions. The world consists of $n+1$ symmetric countries: a home economy (indexed by 0) and n foreign destinations $d \in \{1, 2, \ldots, n\}$. In each destination there is a representative consumer, and firms can choose whether to serve that market. To export to destination d, a firm must incur two standard trade frictions: a destination-specific fixed cost f_x and an iceberg trade cost $\tau > 1$. These frictions govern both the extensive margin (whether the firm exports) and the intensive margin (how much it ships conditional on exporting).

Demand in each destination contains an idiosyncratic component z_d, drawn independently across markets from the same distribution as in autarky. This term captures destination-specific tastes, match quality, or local demand conditions that firms do not observe perfectly when making export decisions. AI adoption improves firms' ability to forecast destination demand by reducing the uncertainty surrounding z_d. As a result, AI affects which markets firms enter and how they allocate sales across destinations once they do export.

Finally, to mirror the empirical regularity that AI adoption is more prevalent among exporters, and that exporters are typically larger and more productive, we impose a sufficiently large fixed export cost f_x. This delivers a sorting pattern in which adoption is more widespread than exporting: some firms find it profitable to adopt AI even if they do not export, while only the most productive firms both adopt and serve foreign markets. We state the condition ensuring this ordering next.

Assumption 3 *The fixed cost of exporting satisfies the following inequality:*

$$\tau \left(\frac{f_x}{\mathbb{E}_z[z^{\sigma-1}]} \right)^{\frac{1}{\sigma-1}} > \left(\frac{f_{AI}}{\mathbb{E}_z[z^{\sigma-1}] - \left(\mathbb{E}_z[z^{\frac{\sigma-1}{\sigma}}]\right)^{\sigma}} \right)^{\frac{1}{\sigma-1}}.$$

This assumption ensures that the marginal firm that adopts AI remains less productive than the marginal exporter, resulting in a nested structure of technology adoption and export behavior.

The presence of fixed export costs and an AI adoption cost implies a simple sorting pattern: all exporters adopt AI, while some non-exporters do not.[16] With countries assumed to be symmetric, wages are identical across markets and normalized to one. For exporters

[16] Appendix B.3.2 derives the cutoff structure: exporting raises the return to AI by improving destination-level quantity choices, so firms that export adopt AI, while less productive firms may stay domestic and not adopt.

that adopt AI, the optimal quantity allocated to each foreign destination d, characterized by consumer preference parameter z_d, and the corresponding price are given by:

$$q^{AI}(s, z_d|e) = \left[\frac{\sigma-1}{\tau\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} (z_d P)^{\frac{\sigma-1}{\sigma}}\right]^{\sigma}, \quad p^{AI}(s, z_d|e) = \frac{\tau\sigma}{\sigma-1} s^{\alpha}. \tag{5}$$

Through AI, exporters can now make more informed production decisions, better aligning output with demand conditions in each foreign market. The pricing rule remains a constant markup over marginal cost, adjusted for the iceberg trade cost τ.

Following the same logic and derivations as in the closed economy case—and using equation (5)—we derive the expected profits of domestic firms with ($g = a$) and without AI ($g = p$), as well as exporters using AI ($g = e$). These derivations are detailed in Appendix B.3.2, and the resulting expected profits are summarized below:

$$\mathbb{E}_z\Big(\pi(s, z|g)\Big) = \begin{cases} \frac{\left[\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right]^{\sigma}}{\sigma^{\sigma}} \Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1} - f, & g = p \\ \frac{\mathbb{E}_z\left(z^{\sigma-1}\right)}{\sigma^{\sigma}} \Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1} - f - f_{AI}, & g = a \\ \frac{(1+n\tau^{1-\sigma})\mathbb{E}_z\left(z^{\sigma-1}\right)}{\sigma^{\sigma}} \Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1} - f - f_{AI} - nf_x, & g = e \end{cases}$$

With fixed costs of production, AI adoption, and exporting, the open-economy equilibrium features three quality cutoffs that govern firm sorting. Firms with $s < s_p^{AI}$ exit; firms with $s \in [s_p^{AI}, s_a^{AI})$ produce domestically without AI; firms with $s \in [s_a^{AI}, s_e^{AI})$ adopt AI but serve only the home market; and firms with $s \geq s_e^{AI}$ both adopt AI and export. As shown in Appendix B.3.2, only sufficiently high-quality firms can cover the additional fixed export cost, so exporting occurs only above the cutoff s_e^{AI}, given by:

$$s_e^{AI} = \tau^{\frac{1}{1-\alpha}} s_p^{AI} \left(\frac{f_x \left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}}{f\left[\mathbb{E}_z(z^{\sigma-1})\right]}\right)^{\frac{1}{(\sigma-1)(1-\alpha)}} > s_a^{AI}.$$

The inequality follows from Assumption 3 and implies a strict sorting hierarchy: $s_p^{AI} < s_a^{AI} < s_e^{AI}$. Intuitively, higher-quality firms are better equipped to absorb the fixed costs associated with both AI adoption and exporting. As a result, they are more likely to upgrade technologies and access foreign markets. This sorting mechanism is illustrated in Figure 1, which depicts the distribution of firms along the quality spectrum. Firms with quality below s_p^{AI} exit the market; those with $s_p^{AI} < s < s_a^{AI}$ remain active but do not adopt AI; firms

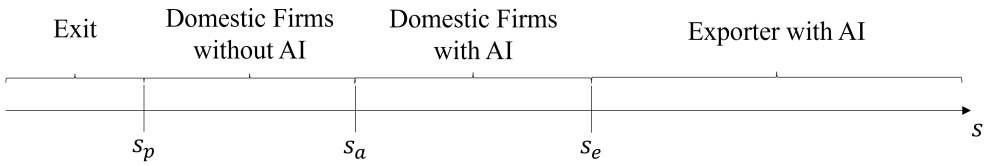


Figure 1: The Sorting Pattern when AI is Available in Open Economy

with $s_a^{AI} < s < s_e^{AI}$ adopt AI for domestic production; and those with $s > s_e^{AI}$ adopt AI and export. The corresponding profit profiles across these thresholds are shown in Figure 2. The free entry condition ensures that expected profits of potential entrants equal the entry cost, given by $\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z(\pi^{AI}(s,z))dF(s) = f_e$, which is solved in Appendix B.3.2. As in previous sections, equilibrium entry is pinned down by the labor market clearing condition. Finally, we prove the existence and uniqueness of equilibrium in this setting in Appendix B.3.3.

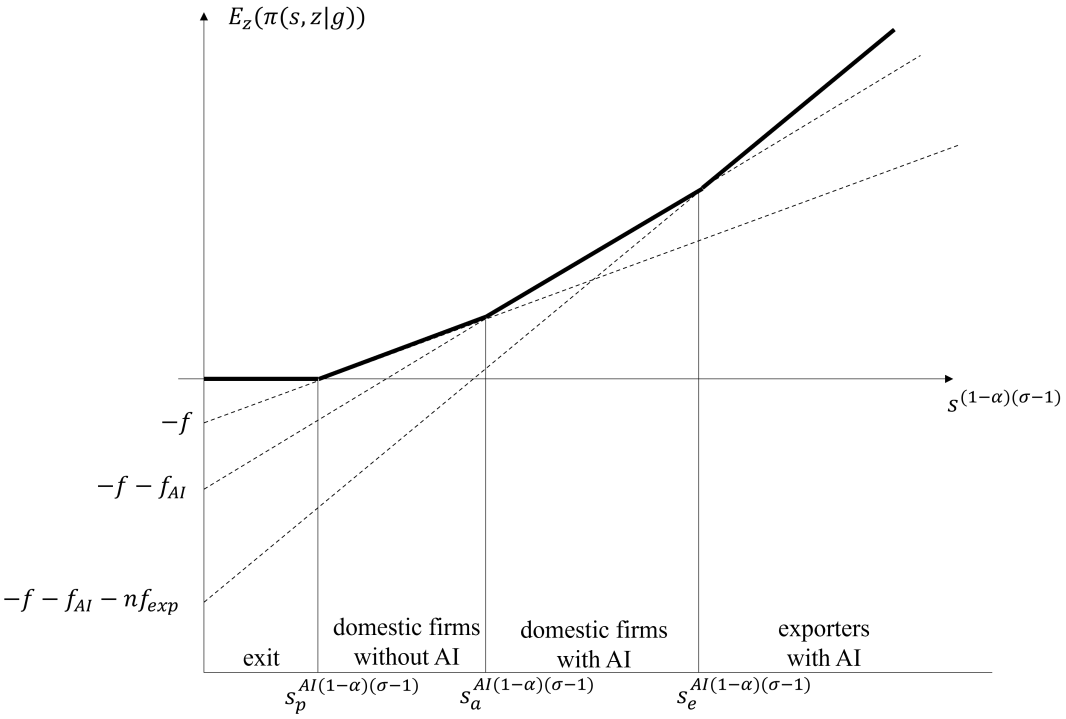


Figure 2: Relationship between Expected Profit and Quality

Impact of AI Adoption on Exporters Having fully described the model, we now examine how the introduction of AI technology alters trade patterns in an open economy. Relative to the baseline equilibrium without AI—where firms face information frictions regarding consumer demand across destinations—AI enables firms to allocate output more efficiently across markets. This improved precision raises expected revenues, prompting

entry by higher-quality firms and shifting the composition of producers toward the more productive end of the distribution. As a result, the quality cutoff for active firms increases, reflecting a stronger selection effect.

More formally, we demonstrate in Appendix B.3.5 that, relative to the equilibrium without AI, the introduction of AI leads to the following general equilibrium outcomes:

Proposition 2 *In an open economy, the introduction of AI technology—compared to the case without AI—leads to the following equilibrium outcomes:*

i) The quality cutoff for active production increases (i.e., $s_p^{AI} > s_p$), reflecting stronger selection into the market;

ii) For a given product, (exporting) firms allocate more quantity to high-demand destinations and less to low-demand ones. Prices adjust accordingly—falling where demand is strong and rising where it is weak—resulting in lower price dispersion but greater quantity dispersion across foreign markets;

iii) The aggregate price index declines, indicating a welfare gain.

Notably, Proposition 2 highlights that the ability to anticipate market-specific demand enables firms to reduce price dispersion across destinations. With more accurate information, firms no longer depend on reactive price adjustments to clear unforeseen excess supply or unmet demand. Instead, they adjust output preemptively: prices fall in markets with strong demand and rise where demand is weak, resulting in more uniform pricing across destinations. By contrast, quantity dispersion increases as firms tailor output more precisely to local preferences.

These adjustments improve allocative efficiency. Reallocation of market share toward higher-quality firms—combined with more targeted output decisions—lowers the aggregate price index. Under free entry, AI adoption amplifies these effects by reinforcing firm selection and reducing resource misallocation, ultimately delivering a welfare gain. In the next section, we empirically test these predictions and examine the mechanisms through which AI shapes international trade.

4 Empirical Method

Guided by the model' s three mechanisms—tighter pricing, finer quantity allocation, and selection into export markets—we take the predictions to the data. Our empirical proxy for adoption is firms' AI-related sales vacancies (AI-Sales), which capture the uptake of

predictive, non-production tools. We examine four export margins: (i) entry and exit across destinations, (ii) changes in product scope, and (iii–iv) the dispersion of export prices and quantities across destinations. Each specification maps a distinct theoretical channel into a testable regression using firm-level evidence.

Identification relies on within-firm variation across products. Although AI-Sales is a firm-wide capability, it should be deployed more intensively where the firm has richer pre-adoption transaction histories and larger stakes. We therefore compare post-adoption changes for a firm' s data-rich (core) products relative to its data-poor (peripheral) products, and difference out the same core–periphery contrast among non-adopting firms. This triple-difference design absorbs firm-wide shocks and common product–time effects, so identification comes from whether adoption generates a larger within-firm shift precisely for products with greater scope for AI learning. For product-level outcomes, we analyze the relationship between a firm's AI-related hiring and its export behavior by estimating the following regression model:

$$y_{ift} = \alpha + \beta AI^{S}_{ift} + \delta_{if} + \delta_{ft} + \delta_{it} + \varepsilon_{ift} \tag{6}$$

where i, f, and t index products (HS-8), firms, and years. The outcome y_{ift} includes the export margins emphasized by the model—such as entry and exit, changes in product scope, and measures of cross-destination dispersion in export prices and quantities. The regressor AI^{S}_{ft} measures firm f's intensity of AI-Sales hiring in year t, which we interpret as a proxy for adoption of predictive AI in sales and market-facing activities. Section 4.1 describes the construction of this measure in detail.

To mitigate potential confounding factors and ensure consistent estimation of the parameters, we include a rich set of fixed effects. Firm–product fixed effects (δ_{if}) absorb all time-invariant determinants of a firm's performance in a given product—such as baseline quality, product-specific marginal costs, or persistent match quality. Firm–year fixed effects (δ_{ft}) capture any firm-wide shocks that vary over time, including changes in productivity, management, financing conditions, or overall export strategy. Importantly, δ_{ft} also helps address concerns about endogenous AI adoption: any unobserved shock that both raises a firm's propensity to hire AI workers and affects its export outcomes is differenced out as long as it operates at the firm-year level. Product–year fixed effects (δ_{it}) absorb common shocks to a product in a given year, such as global demand shifts, commodity-price movements, or changes in product-specific trade policy. Standard errors are clustered at the industry level to allow for arbitrary correlation in outcomes and adoption within industries over time.[17]

[17]For firm-level outcomes, such as export participation and product scope, we estimate $y_{f(cj)t} = \alpha + \beta AI^{S}_{ft} + \delta_f + \delta_{ct} + \delta_{jt} + \varepsilon_{f(cj)t}$. These specifications include firm fixed effects (δ_f) to absorb time-invariant firm heterogeneity, industry–year fixed effects (δ_{jt}) to control for sectoral shocks, and city–year fixed effects

4.1 Measurement

4.1.1 Output and Price Adjustment

The adoption of AI—proxied by firm-level AI-Sales hiring—may enable exporters to better understand consumer demand and optimize their export strategies. As outlined in Proposition 2, AI adoption improves firms' ability to identify destination-specific preferences for a given product. This leads to a more targeted allocation of output across markets, resulting in greater dispersion in quantities sold and, conversely, more uniform pricing, as firms no longer need to rely on price adjustments to manage inventory imbalances under uncertainty.

To test this hypothesis, we construct two dispersion measures—one for quantity and one for price—defined as:

$$Var(q)_{ift} = \left(\frac{\sum_{d \in D_{ift}} \left(q_{ifdt} - \overline{q}_{ift} \right)^2}{|D_{ift}|} \right), \quad Var(p)_{ift} = \left(\frac{\sum_{d \in D_{ift}} \left(p_{ifdt} - \overline{p}_{ift} \right)^2}{|D_{ift}|} \right) \tag{7}$$

where q_{ifdt} and p_{ifdt} denote, respectively, the quantity and price of product i exported by firm f to destination d in year t. The terms $\overline{q}_{ift}$ and $\overline{p}_{ift}$ represent the average quantity and price across all destinations D_{ift} served by that firm-product-year combination.

The resulting variables $Var(q)_{ift}$ and $Var(p)_{ift}$ capture the cross-destination variance in quantities and prices, respectively. To ensure these measures are meaningful, we restrict the sample to firm-product pairs that are exported to at least two destinations in a given year. A higher value of $Var(q)_{ift}$ indicates greater variation in quantities allocated across destinations—suggesting more effective matching between supply and local demand. A lower value of $Var(p)_{ift}$ reflects reduced price dispersion—suggesting that firms rely less on price variation to cope with uncertain demand. In our regressions, we take the natural logarithm of these dispersion measures. Theory predicts AI lowers price dispersion but raises quantity dispersion across destinations.

In addition to analyzing changes in price and quantity dispersion, we also examine how a firm's adoption of AI affects its decisions to enter or exit export markets at the product level. To do so, we construct measures of market entry and exit by comparing each firm-product's set of export destinations in year t to the baseline year, 2013. A destination is classified as a market entry if product i is exported to destination d in year t, but not in 2013; similarly, a destination is counted as a market exit if it was served in 2013 but not in

(δ_{ct}) to account for local policies and macroeconomic conditions. Robust standard errors are clustered at the industry level.

year t. We define N_{ift}^{entry} and N_{ift}^{exit} as the number of market entries and exits, respectively, for each firm-product-year observation. At the firm level, we construct additional measures to capture export participation and product scope. The binary variable $\mathbb{I}_{ft}^{x}$ equals one if firm f exports in year t, and zero otherwise, capturing the extensive margin of export activity. To quantify the firm's product scope, we define N_{ft}^{pdt} as the number of distinct HS8-level product categories exported by firm f in year t.

4.1.2 Exposure to AI Technology

To study how AI adoption in non-production activities shapes exporters' outcomes, we construct a product-level measure of exposure to AI–Sales technology. As described in Section 2, we first identify whether a firm posts at least one vacancy that both lists an AI-related skill and references sales- or market-facing tasks in the job description. We summarize this information with a firm-year indicator AI_{ft}^{S}, equal to one if firm f has posted an AI–Sales vacancy by year t, and zero otherwise.[18] To reflect the idea that adoption builds a persistent capability, once a firm enters AI–Sales hiring we set $AI_{ft}^{S} = 1$ in all subsequent years. We then translate this firm-wide adoption measure into product-level exposure using the importance of each product in the firm's export portfolio. Intuitively, AI–Sales capability is likely to be deployed first, and most intensively, on a firm's core export lines. We therefore define exposure for product i within firm f as:

$$AI_{ift}^{S} = Share_{if,2013} \times AI_{ft}^{S}, \text{ with } Share_{if,2013} = \frac{x_{if,2013}}{\sum_{i \in I_{f,2013}} x_{if,2013}} \tag{8}$$

where $Share_{if,2013}$ is product i's share in firm f's total export value in 2013, $x_{if,2013}$ is export value, and $I_{f,2013}$ is the set of products exported by firm f in 2013.[19] In effect, AI_{ift}^{S} treats AI–Sales as a firm-wide capability whose relevance is greater for products that are more central to the firm's export mix.[20] Allowing exposure to vary across products within a firm is motivated by three considerations. First, firms typically concentrate managerial attention and complementary investments on their core products, where the returns to improved forecasting and targeting are likely to be highest.[21] Second, high-volume products generate

[18] As a robustness check, we also construct a continuous measure that captures the intensity of AI–Sales hiring; results are reported in the robustness section.

[19] We use baseline export shares rather than levels because export values vary widely across firms; shares provide a scale-free measure of a product's importance within the firm. We freeze product weights at their 2013 baseline values so that variation in AI_{ift}^{S} is driven by adoption timing, not by endogenous changes in product mix following adoption.

[20] For example, if product A accounts for 20% of a firm' s exports in 2013 while product B accounts for 2%, then after adoption we set $AI_{A,ft}^{S} = 0.20$ and $AI_{B,ft}^{S} = 0.02$; before adoption, both values are zero.

[21] Bernard et al. (2007) show that exporters disproportionately emphasize products tied to their core competencies, consistent with strategic resource allocation toward higher-value lines.

more transactions and richer feedback, which expands the scope for data-driven learning in pricing and demand prediction.[22] Third, weighting by baseline export shares mitigates measurement error by down-weighting marginal products that are shipped infrequently.

Relative to the alternative assumption that all products benefit equally from AI adoption, this construction yields a more targeted proxy for effective exposure within the firm. Moreover, fixing product weights at their 2013 baseline values ensures that the estimated effects are driven by changes in AI adoption rather than endogenous adjustments in product mix. In the robustness section, we also report results from a simpler specification that classifies products as core versus non-core without using the exact share information; the conclusions are unchanged.

4.2 Endogeneity and Instrumental Variable Strategy

Two identification concerns merit emphasis. First, *core* products may follow different pricing and allocation patterns than peripheral products for reasons unrelated to AI. Our exposure measure uses baseline export shares, fixed in 2013, so it is not mechanically affected by subsequent export performance. All specifications also include firm–product fixed effects δ_{if}, which absorb time-invariant firm–product attributes such as persistent comparative advantage, product-specific markups, or long-run customer relationships. Second, AI adoption itself may respond to unobserved shocks. Firms facing rising cross-destination demand volatility may be more likely to adopt AI to improve forecasting and targeting. Likewise, broader upgrading—improvements in management, IT systems, or production efficiency—could raise the propensity to hire AI-related workers and independently improve export allocation. Our baseline design mitigates this concern by including firm–year fixed effects, so identification comes from within-firm reallocations across products within the same year, holding constant any firm-wide shocks that could jointly affect AI hiring and export outcomes.

To further address this concern, we develop an IV strategy that exploits plausibly exogenous variation in local access to AI-skilled labor. The instrument builds on a distinctive feature of China's higher-education system: university admissions are governed by centrally planned enrollment quotas. Each year, the Ministry of Education (MoE), together with provincial authorities and universities, approves field-by-region intake targets. These quotas are shaped by national development priorities and medium- to long-run workforce planning, rather than the short-run hiring needs of particular firms.[23] Using these approved admis-

[22] Dickstein and Morales (2018) emphasize that larger sales volumes provide more information about market conditions, strengthening the link between scale and learning.

[23] In China, admissions plans are coordinated by the MoE and provincial education authorities. Universities propose intake targets, but quotas by major are subject to approval and are guided by policy priorities and long-run labor market planning. See, for example, the MoE's enrollment guidelines: http://www.moe.gov.

sions plans, we construct a firm-level measure of potential access to AI talent based on the geographic distribution of AI-related enrollment. Intuitively, larger nearby quotas raise the local supply of relevant skills and reduce firms' cost of recruiting AI workers, but are unlikely to move contemporaneous export outcomes except through AI adoption. Formally, we define:

$$AI_{ft}^{IV} = \sum_{u \in U} \frac{1}{dist_{fu}} \times AIQuota_{u,t-1}, \tag{9}$$

where $dist_{fu}$ is the distance between firm f and university u, U is the set of universities, and $AIQuota_{u,t-1}$ is the number of students admitted in year $t-1$ to AI-related majors (e.g., computer science, data science, and artificial intelligence) at university u, as recorded in the MoE's officially approved enrollment plans.

Using lagged quotas helps address endogeneity by shifting firms' access to AI talent for reasons external to the firm, rather than reflecting endogenous adoption in response to unobserved shocks. The IV measures a firm's distance-weighted exposure to centrally approved AI enrollment quotas at nearby universities. Larger nearby quotas increase the local flow of AI-trained graduates and lower recruiting frictions, making AI hiring more likely even absent firm-specific demand shocks or broad upgrading. Because quotas are set in advance as part of national workforce planning and enter with a lag, this variation is predetermined with respect to contemporaneous firm-level shocks, mitigating reverse causality.[24] In the IV analysis, we instrument product-level sales-AI exposure by replacing the firm adoption component in equation (8) with AI_{ft}^{IV}, yielding AI_{ift}^{IV} as the instrument, and we also control for non-sales AI hiring and broader data-technology adoption so the identifying variation is not simply general digital upgrading.

5 Empirical Results

5.1 Baseline Results

We begin with the baseline specification in equation (6), which relates AI adoption to cross-destination price and quantity dispersion for a given firm–product. Panel A of Table 3 reports the results. Columns (1) and (3) examine price and quantity dispersion, respectively. Because dispersion can mechanically vary with the number of destinations served, Columns (2) and (4) additionally control for the log number of export destinations.

cn/jyb_xwfb/moe_2082/zl_2016n/2016_zl30/201606/t20160606_248332.html.

[24] Appendix Figure A.3 plots the geographic distribution of AI enrollment capacity (green dots) and the density of AI-related vacancies by prefecture (shading), showing substantial overlap between talent supply and AI hiring.

Two robust patterns emerge. First, after a firm begins hiring for AI–Sales roles, price dispersion across destinations falls for products that are more central to the firm 's pre-adoption export mix (Columns (1)–(2)). Second, for the same products, quantity dispersion rises (Columns (3)–(4)). Controlling for the number of destinations leaves both the magnitudes and statistical significance essentially unchanged. For a product with the mean export share (6.2%), the estimates imply that AI–Sales adoption is associated with a 14% decline in cross-destination price dispersion and a 2% rise in cross-destination quantity dispersion for the representative firm–product.[25] These findings match the model's predictions: by reducing demand-side information frictions, AI adoption allows firms to reallocate quantities toward destinations with stronger demand while harmonizing prices across markets.

Table 3: Impact of AI Adoption on Price and Quantity Dispersion

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(q)_{ift}$	(4) $\ln Var(q)_{ift}$
Panel A: Price and Quantity Dispersion Across Destinations				
AI^S_{ift}	-2.187***	-2.202***	0.309**	0.304**
	(0.497)	(0.526)	(0.137)	(0.129)
$\ln(\#\text{ of Destinations})_{ift}$		0.760***		0.413***
		(0.014)		(0.018)
Observations	106,510	106,510	105,210	105,210
Adjusted R^2	0.788	0.793	0.902	0.903
Panel B: Price and Quantity Dispersion Across All Transactions				
AI^S_{ift}	-2.508***	-2.490***	0.266**	0.278**
	(0.514)	(0.541)	(0.122)	(0.121)
$\ln(\#\text{ of Transactions})_{ift}$		0.794***		0.393***
		(0.011)		(0.010)
Observations	113,250	113,250	111,820	111,820
Adjusted R^2	0.778	0.783	0.900	0.901
Fixed Effects:				
Firm-HS 8-digit	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y

Note: Panel A reports dispersion in prices and quantities across destinations within a firm–product–year, calculated as the variance in average unit prices and quantities sold across all destination markets that the firm serves with that product in a given year, as defined in equation (7). Panel B measures dispersion across individual transactions within each firm–product–year, using the full set of shipment-level trade flows, including repeated transactions to the same destination. $\ln(\#\text{ of Destinations})_{ift}$ and $\ln(\#\text{ of Transactions})_{ift}$ denote the log number of destinations and total transactions, respectively, for firm f selling product i in year t. Robust standard errors are clustered at the industry level and reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

[25] These magnitudes are obtained by multiplying the estimated coefficient on AI^S_{ift} by the mean of $Share_{if,2013}$ (6.2%).

Transaction-Level Robustness. In the baseline estimations, we aggregate export transactions at the product-destination-year level before computing dispersion measures. As a robustness check, we construct dispersion measures at the transaction level by calculating the variance in prices and quantities across all transactions for a given firm-product-year.[26]

Panel B of Table 3 presents the regression results using these transaction-level dispersion measures. Consistent with the baseline findings, we observe that price variance is negatively associated with a product's AI-Sales exposure, while quantity variance is positively associated. These results reinforce our interpretation that AI-Sales adoption enhances firms' ability to coordinate pricing across destinations while reallocating output toward markets with stronger demand. he IV results in Table 4 further strengthen this conclusion. Column (1) shows a strong first stage: the instrument is highly predictive of AI–Sales use intensity. Columns (2)–(5) report second-stage estimates for price and quantity deviations that mirror the OLS findings. The IV coefficients are larger in absolute value, consistent with attenuation in OLS due to measurement error in the firm–product–year AI exposure measure.[27] The parallel results across specifications reinforce our main conclusions.

Other Export Performance. To further probe how AI adoption reshapes exporters 'choices, we examine destination-market adjustment along the extensive margin at the firm–product level. Column (1) of Table 5 shows that after AI adoption, products with larger baseline export shares enter more new destinations (relative to 2013). Column (2) shows that these same products also exit more destinations—those served in 2013 but no longer served in year t. The joint rise in entry and exit indicates active market re-optimization: firms add destinations with stronger demand prospects and drop weaker ones. This pattern also helps rule out a simple simultaneity story in which the AI–Sales measure merely proxies for broad firm expansion or generalized upgrading. Such a story would predict more entry but fewer exits, especially for core products. Instead, exits increase alongside entries for precisely those high-share products, pointing to deliberate demand-driven market selection rather than mechanical growth.

We then move to firm-level extensive-margin outcomes. Columns (3)–(6) estimate an

[26]Specifically, we define transaction-level dispersion as $Var(q)_{ift} = \sum_{e \in E_{ift}} \left(q_{ifet} - \overline{q}_{ift}\right)^2 / |E_{ift}|$ and $Var(p)_{ift} = \sum_{e \in E_{ift}} \left(p_{ifet} - \overline{p}_{ift}\right)^2 / |E_{ift}|$, where q_{ifet} and p_{ifet} denote the quantity and price of product i sold by firm f in transaction e in year t, and $\overline{q}_{ift}$, $\overline{p}_{ift}$ are the corresponding means across all transactions E_{ift}.

[27]With classical measurement error, the OLS probability limit satisfies plim $\hat{\beta}^{OLS} = \beta \cdot \frac{\text{Var}(AI^S ift)}{\text{Var}(AI^S ift) + \text{Var}(\nu_{ift})}$, where ν_{ift} denotes measurement noise. Such noise can arise when mapping AI application intensity to products using sales weights and when AI use is observed imperfectly (e.g., via a dummy indicator), which attenuates OLS estimates toward zero.

Table 4: AI Adoption on Cross-Market Price and Quantity Dispersion: IV Estimates

Dep var.	(1) First Stage AI^S_{ift}	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(p)_{ift}$	(4) $\ln Var(q)_{ift}$	(5) $\ln Var(q)_{ift}$
Panel A: Price and Quantity Dispersion Across Destinations					
AI^{IV}_{ift}	0.108***				
	(0.011)				
AI^S_{ift}		-31.79***	-31.31***	1.412**	1.676**
		(4.629)	(4.526)	(0.597)	(0.653)
$\ln(\#\text{ of Destinations})_{ift}$			0.756***		0.413***
			(0.015)		(0.018)
Observations	106,500	106,500	106,500	105,200	105,200
F Statistics	99.31				
Panel B: Price and Quantity Dispersion Across All Transactions					
AI^{IV}_{ift}	0.107***				
	(0.011)				
AI^S_{ift}		-37.04***	-36.36***	1.706***	2.057***
		(5.629)	(5.511)	(0.567)	(0.615)
$\ln(\#\text{ of Transactions})_{ift}$			0.784***		0.394***
			(0.013)		(0.010)
Observations	113,240	113,240	113,240	111,810	111,810
F Statistics	97.62				
Fixed Effects:					
Firm-HS 8-digit	Y	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y	Y

Note: Panel A reports IV results on the dispersion of prices and quantities across destinations for each firm–product–year, measured as the variance in average unit prices and quantities sold across all destination markets served with that product, as defined in equation (7). Panel B reports dispersion measured across individual transactions within each firm–product–year, using the full set of shipment-level records, including repeated transactions to the same destination. The variables $\ln(\#\text{ of Destinations})_{ift}$ and $\ln(\#\text{ of Transactions})_{ift}$ denote the log number of export destinations and total transactions, respectively, for firm f exporting product i in year t. We control for AI not identified as for sales purpose (AI^{NS}_{ift}) and other data technology ($Data_{ift}$). Robust standard errors, clustered at the industry level, are reported in parentheses. *** p <0.01, ** p <0.05, * p <0.1.

Table 5: AI Adoption and Export Entry, Exit, and Scope

Dep var.	(1) N_{ift}^{entry}	(2) N_{ift}^{exit}	(3) $\mathbb{I}_{ft}^{x}$	(4) $\mathbb{I}_{ft}^{x}$	(5) N_{ft}^{pdt}	(6) N_{ft}^{pdt}
AI_{ift}^{S}	0.971**	0.301**				
	(0.358)	(0.116)				
AI_{ft}^{S}			0.019***	0.019***	0.062**	0.069**
			(0.002)	(0.002)	(0.025)	(0.027)
Observations	214,742	214,742	133,883	131,704	29,422	28,642
Adjusted R^2	0.815	0.896	0.739	0.736	0.901	0.904
Fixed Effects:						
Firm-HS 8-digit	Y	Y				
Firm-Year	Y	Y				
HS 8-digit-Year	Y	Y				
Firm			Y	Y	Y	Y
Industry-Year			Y	-	Y	-
City-Year			Y	-	Y	-
Industry-City-Year			-	Y	-	Y

Note: Variables N_{ift}^{entry} and N_{ift}^{exit} denote the number of destination markets newly entered and exited, respectively, for each firm-product-year observation. At the firm level, the binary indicator $\mathbb{I}_{ft}^{x}$ equals one if firm f engages in exporting in year t, and zero otherwise. N_{ft}^{pdt} represents the number of distinct HS 8-digit product categories exported by firm f in year t. Columns (1) and (2) report results from firm-product-level regressions, where the dependent variables are the number of newly entered and exited export markets, relative to each product's 2013 market portfolio. Columns (3) through (6) present firm-level regressions. The dependent variable in Columns (3) and (4) is export status $\mathbb{I}_{ft}^{x}$, while in Columns (5) and (6), it is the product scope N_{ft}^{pdt}, conditional on exporting. Columns (3) and (5) include firm, industry-year, and city-year fixed effects, while Columns (4) and (6) replace industry-year and city-year fixed effects with industry-city-year fixed effects. Robust standard errors, clustered at the industry level, are reported in parentheses. *** p <0.01, ** p <0.05, * p <0.1.

analogous specification at the firm level, with the dependent variable capturing export participation in Columns (3)–(4) and product scope (the number of HS8 varieties exported, conditional on exporting) in Columns (5)–(6). The key regressor is AI^S_{ft}, an indicator for whether firm f hires for AI–Sales roles in year t. We include firm fixed effects, industry–year fixed effects, and city–year fixed effects; Columns (4) and (6) further absorb industry×city×year shocks. Standard errors are clustered at the industry level. Across specifications, AI adoption raises the probability of exporting by 1.9 percentage points. Relative to a baseline export participation rate of 23%, this corresponds to an 8.3% increase. Conditional on exporting, AI adopters also expand product scope, exporting about 0.06–0.07 additional HS8 varieties.

Finally, although we emphasize allocation across destinations, these reallocations can also move average export prices and quantities. The sign is a priori ambiguous because averages reflect two distinct allocation margins: how a firm reallocates sales across destinations it already serves, and which destinations it chooses to serve in the first place.[28] Appendix D.2 reports the corresponding evidence: average prices decline and average quantities increase modestly, with the latter estimated less precisely.

5.2 Heterogeneous Effects of Sales-Oriented AI Adoption

Product Differentiation. Panel A of Table 6 examines whether the effects of AI–Sales vary with product differentiation. The model implies larger gains where demand is more information-intensive—that is, when consumer tastes, design, and market-specific appeal play a central role. We therefore expect stronger effects for differentiated goods than for homogeneous ones. We implement this test in two steps. First, we use the UN Broad Economic Categories (BEC) classification to separate intermediate from non-intermediate products.[29] Because non-intermediate goods are purchased by end consumers and are more likely to be shaped by taste and variety considerations, they are the natural setting in which sales-oriented AI should matter most; intermediate inputs, in contrast, are typically procured on specifications and are less sensitive to such demand-side variation. We therefore focus on the non-intermediate sample. Second, within this sample we interact AI^S_{ift} with $Diff_i$, an indicator for whether product i is classified as differentiated under BEC. The interaction coefficients are statistically significant and have the predicted signs, indicating

[28] On the intensive margin, improved demand forecasts shift quantities toward stronger markets and away from weaker ones, with prices adjusting in the opposite direction. On the extensive margin, AI is associated with market turnover—entry into high-demand destinations and exit from low-demand ones where prices would otherwise be relatively low.

[29] Product groupings are based on the Broad Economic Categories (BEC) classification and mapped to HS 8-digit codes using the official correspondence provided by the United Nations Statistics Division (https://unstats.un.org/unsd/classifications/econ/).

that the effects of AI–Sales are indeed amplified for differentiated products.

Table 6: Heterogeneous Effects of Sales-Oriented AI on Price and Quantity Dispersion

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(q)_{ift}$	(4) $\ln Var(q)_{ift}$	(5) $\ln Var(p)_{ift}$	(6) $\ln Var(p)_{ift}$	(7) $\ln Var(q)_{ift}$	(8) $\ln Var(q)_{ift}$
Panel A: By Product Differentiation								
$AI^S_{ift} \times Differentiated_i$	-1.453***	-1.136**	2.096***	2.241***	-1.225***	-0.854***	2.132***	2.288***
	(0.316)	(0.261)	(0.646)	(0.690)	(0.405)	(0.342)	(0.631)	(0.671)
$\ln(\#\text{ of Destinations})_{ift}$		Y		Y				
$\ln(\#\text{ of Transactions})_{ift}$						Y		Y
Observations	42,976	42,976	42,096	42,096	45,056	45,056	44,088	44,088
Adjusted R^2	0.800	0.805	0.901	0.902	0.792	0.797	0.899	0.900
Panel B: By Firm Size								
$AI^S_{ift} \times Large_f$	-1.235	-1.210	0.950***	0.968***	-0.957	-0.857	1.125***	1.180***
	(0.732)	(0.740)	(0.287)	(0.270)	(0.637)	(0.652)	(0.334)	(0.315)
$\ln(\#\text{ of Destinations})_{ift}$		Y		Y				
$\ln(\#\text{ of Transactions})_{ift}$						Y		Y
Observations	106,510	106,510	105,210	105,210	113,250	113,250	111,820	111,820
Adjusted R^2	0.788	0.793	0.902	0.903	0.778	0.783	0.900	0.901
Fixed Effects:								
Firm-HS 8-digit	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y	Y	Y	Y	Y
Panel C: By Destination Market								
$AI^S_{ift} \times Developed_g$	-0.0439	-0.153	0.656**	0.605**	-0.255	-0.334	0.791***	0.761***
	(0.861)	(0.808)	(0.294)	(0.321)	(0.585)	(0.543)	(0.263)	(0.290)
$\ln(\#\text{ of Destinations})_{ifgt}$		Y		Y				
$\ln(\#\text{ of Transactions})_{ifgt}$						Y		Y
Observations	116,996	116,996	115,620	115,620	127,272	127,272	125,722	125,722
Adjusted R^2	0.774	0.779	0.891	0.892	0.765	0.770	0.889	0.890
Fixed Effects:								
Firm-HS 8-digit-Destination Group	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Destination Group-Year	Y	Y	Y	Y	Y	Y	Y	Y
HS 8-digit-Destination Group-Year	Y	Y	Y	Y	Y	Y	Y	Y

Note: Panel A examines heterogeneity by product differentiation. Panel B examines heterogeneity by firm size. Panel C examines heterogeneity by destination type. Columns (1)–(4) report price and quantity dispersion across destinations, and Columns (5)–(8) across transactions. Dispersion is calculated as the variance in average unit prices and quantities defined in equation (7). $\ln(\#\text{ of Destinations})_{ift}$ denotes the log number of export destinations for firm f selling product i in year t, and $\ln(\#\text{ of Transactions})_{ift}$ denotes the log number of export transactions for firm f selling product i in year t. In panel C, $\ln(\#\text{ of Destinations})_{ifgt}$ and $\ln(\#\text{ of Transactions})_{ifgt}$ denotes the log number of export destinations and transactions within destination group g for firm f selling product i in year t respectively. Robust standard errors, clustered at the industry level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Firm Size. We next examine heterogeneity by firm size. We classify firms as *large* if their employment is in the top 30% of the firm distribution in our sample, and as SMEs otherwise, and interact AI^S_{ift} with an indicator $Large_f$. Panel B of Table 6 reports the results: Columns (1)–(4) use dispersion across destinations, and Columns (5)–(8) use dispersion across transactions. The effects are stronger for large firms. While the differential effect on price dispersion is imprecisely estimated, the corresponding difference for quantity dispersion

is positive and statistically significant at the 1% level. This pattern is consistent with the idea that AI is most productive when it can be trained on rich, high-frequency information—a condition more likely to hold for larger firms with greater scale and data. In this sense, AI adoption can amplify pre-existing capability differences, echoing the learning advantage of large firms emphasized by Dickstein and Morales (2018).

Market Income and Consumer Preferences. We further explore heterogeneity across destination markets by splitting export destinations into developed and developing economies. The dispersion outcomes are defined as in equation (7), but we now compute price and quantity dispersion separately within each destination group. We then interact AI^S_{ift} with an indicator $Developed_g$ that equals one when dispersion is computed within the set of *OECD-defined developed economies*.[30] The specification absorbs group-specific shocks by including firm–HS8–group fixed effects, firm–group–year fixed effects, and HS8–group–year fixed effects. Panel C of Table 6 reports the results: Columns (1)–(4) use dispersion across destinations, and Columns (5)–(8) use dispersion across transactions. The effects of AI–Sales are substantially stronger within developed markets, especially for quantity dispersion. This pattern is consistent with three complementary forces. First, demand is typically more differentiated in high-income economies, where product attributes such as design, features, and branding play a larger role in purchase decisions, raising the value of precise demand prediction.[31] Second, firms serve developed destinations more often, which expands the data available for learning and improves the performance of data-intensive tools.[32] Third, learning about foreign demand tends to be harder in developing markets, which lowers the returns to fine targeting and can limit the effectiveness of forecasting-based tools.[33] The results support our interpretation that AI–Sales sharpens demand-side information, with the largest effects where preferences are richer, feedback is more abundant, and information barriers are lower.

Overall, we find that sales-oriented AI is most effective in environments where consumer heterogeneity plays a central role. Its impact is especially pronounced for non-intermediate differentiated goods, where products are less standard and end consumers base purchasing decisions on taste, design, and other subjective attributes. In these settings, the ability to interpret and respond to diverse demand signals becomes a strategic advantage—one that

[30] We classify destination markets using the OECD country-income/development grouping and map it to our destination codes.

[31] A large literature documents that higher-income countries consume a wider range of varieties, including imported varieties; see Jackson (1984), Hunter and Markusen (1988), and Simonovska (2015).

[32] Wagner (2018) shows that firms are more likely to export to high-income destinations, consistent with greater exposure and richer feedback from these markets.

[33] Allen (2014) emphasizes that the costs of acquiring and processing foreign market information are higher in poorer destinations.

tends to benefit larger firms disproportionately. With access to richer data and greater analytical capacity, these firms are better equipped to translate AI-driven insights into actionable market strategies. The effectiveness of sales AI also varies systematically across destination markets. In developed economies, where consumers exhibit greater preference diversity and assign higher value to product features, the gains from demand-side learning are more substantial.

6 Robustness Analysis

We perform a variety of sensitivity checks to ensure the robustness of our results. We address outliers and alternative treatments of missing vacancy data, vary how AI exposure and sales-oriented AI are defined, contrast sales-oriented AI with other AI uses, and run placebo tests to rule out spurious trends.

Removing Influential Observations. AI-related hiring is concentrated in China's largest municipalities, where access to talent and complementary digital infrastructure is strongest. To ensure that our results are not driven by a few such locations, we re-estimate the baseline regressions after excluding firms headquartered in Beijing, Tianjin, Shanghai, and Chongqing. Panel A of Table A.4 reports the estimates. The coefficient on AI^S_{ift} remains negative and statistically significant at the 1% level for price dispersion, with only modest attenuation relative to the baseline. In the quantity-dispersion regression, the coefficient remains positive but is less precisely estimated (significant at the 10% level) due to the smaller sample. The point estimate remains close to the main result.

We also verify that the results are not driven by extreme values of the AI–Sales measure. Panel B reports estimates after winsorizing the AI exposure variable at the top and bottom 1% of its distribution. This is relevant because dispersion measures can be sensitive to influential observations, and a small number of firms may undertake unusually intensive AI hiring concentrated in a few flagship export products. The results are unchanged: coefficient magnitudes and significance levels remain close to the baseline, indicating that the findings reflect broad patterns rather than a few outliers.

Addressing Missing Job Vacancies. Firms may not consistently post AI-related job vacancies online each year, and the unavailability of postings could introduce bias if it occurs systematically.[34] To address this concern, we adopt the following modification in our analysis:

[34]For instance, if vacancies are more likely to be posted during periods of rapid growth or expansion, our measure could overstate AI adoption by capturing only the most active periods. Conversely, if firms invest

we treat missing vacancy data as indicative of zero AI-related recruitment activity. This conservative approach assumes that in years without observed postings, firms do not actively recruit AI sales staff, thus aligning the constructed measure more closely with the actual hiring behavior of firms. Panel C of Table A.4 presents the results using this adjusted measure. The estimated coefficients are similar to the baseline results and retain their statistical significance, reinforcing the robustness of the relationship between AI adoption and firms' export reallocation behavior.

Adjusting Product Shares and Core Product Identification. Our main specification interacts a firm's AI-Sales adoption decision, AI^S_{ft}, with each product's export share within the firm in the initial year. However, since export shares can be mechanically influenced by the number of varieties a firm sells, this raises comparability concerns across firms. To address this, we first normalize each product's share by the firm's highest export share, rescaling the core product to one. Using this adjusted measure, we re-estimate the regressions. As shown in Panel D of Table A.4, the coefficients on AI^S_{ift} remain highly significant with the expected signs, although somewhat smaller in magnitude due to the change in scale.

As a further check, we also replace the continuous share measure with a dummy variable that indicates whether a product is the firm's top-selling product. This alternative specification effectively implements a triple-difference framework that compares changes in price and quantity dispersion between core and non-core products before and after AI adoption. The results, reported in Panel E, again show that AI-Sales adoption leads to a larger reduction in price dispersion and a greater increase in quantity dispersion for core products relative to others, reinforcing the interpretation that AI adoption improves firms' ability to allocate output more efficiently across products.

Differentiating Sales and Non-Sales AI. So far, we have focused on AI hiring for non-production tasks, which we label *AI-Sales*. We now broaden the analysis in two directions: (i) overall AI hiring, without distinguishing tasks, and (ii) AI hiring for production-related activities ("non-sales AI"), defined as AI-related vacancies that do not contain any of the non-production keywords.[35] If the mechanism operates primarily through better demand inference, non-sales AI should have markedly weaker effects on our dispersion outcomes. To

in AI capabilities without public postings—such as through internal promotions—missing observations could lead us to understate adoption.

[35]Using the AI keywords in Table A.1, we identify all AI-related vacancies and then exclude those classified as sales AI to construct non-sales AI hiring. Non-sales AI can mechanically compress price dispersion if it lowers marginal costs uniformly across destinations. Because it does not directly improve demand learning, however, it should have little effect on quantity dispersion, which is the key margin of market-specific reallocation.

test this, we construct AI^{NS}_{ift} analogously to AI^{S}_{ift} in equation (8), replacing the adoption indicator with AI^{NS}_{ft}, which equals one once firm f hires non-sales AI workers by year t, and re-estimate the baseline regressions.

Panels A and B of Table A.7 confirm this prediction. Relative to the baseline, non-sales AI has essentially no effect on quantity dispersion and only a small, though negative, association with price dispersion. The null quantity result is especially diagnostic: it indicates that production-side AI does not induce the destination-level reallocations that our mechanism attributes to improved demand information. The modest reduction in price dispersion is consistent with uniform cost savings that compress prices without changing market allocation. Panel B delivers the same message using transaction-level dispersion measures. Panel C then pools all AI hiring, combining sales and non-sales applications. The pooled estimates closely track the baseline pattern, with magnitudes that lie between the sales-only and non-sales results, consistent with sales-oriented applications accounting for most of the economically relevant variation in AI adoption in our setting.[36]

AI–Sales Hiring versus Traditional Sales Hiring. A natural benchmark is conventional sales hiring. Traditional sales teams also collect and process market information, so if our mechanism is fundamentally about easing demand-side information frictions, one would expect similar qualitative patterns for sales hiring more broadly. We therefore compare the effects of AI–Sales hiring with those of hiring in traditional sales occupations. Table A.5 reports the results. Two findings stand out. First, traditional sales hiring is associated with lower price dispersion and higher quantity dispersion across destinations, consistent with the view that sales functions help firms learn about heterogeneous market conditions and reallocate shipments accordingly. Second, the corresponding coefficients for AI–Sales hiring are larger. This gap suggests that AI-enabled sales tools provide a more effective way to extract signal from past transactions and market feedback, improving firms' ability to predict demand and target destinations—particularly in settings where consumer tastes are complex and difficult to summarize with simple heuristics.

Controlling for Broader Data Technology Adoption. A potential concern is that our sales-AI measure may simply proxy for a broader digital capability—firms that invest in data systems and analytics may both hire AI-related workers and adjust export outcomes for reasons unrelated to the demand-learning mechanism. To address this, we construct a sepa-

[36] As an external check, Table A.10 relates AI hiring to firm-level productivity. Non-sales AI is positively associated with productivity, while sales AI is not. This split reinforces our interpretation: the dispersion effects of AI-Sales reflect demand-side learning and reallocation, rather than broad improvements in operating efficiency.

rate keyword list capturing *non-AI* data technologies (Table A.1) and build a corresponding exposure measure, $Data_{ift}$, constructed in the same way as our baseline exposure but based on these general data-technology keywords.[37] We then include $Data_{ift}$ as an additional control in the baseline regressions. Panel D of Table A.7 shows that the coefficient on sales AI remains stable in sign and significance, indicating that the main results are not driven by broader data-technology adoption. In contrast, general data technology is associated with lower price dispersion but has little effect on quantity dispersion.

Panel E strengthens this point by controlling simultaneously for (i) sales AI, (ii) general data technology, and (iii) *other AI hiring outside sales-related roles*. The results continue to show a clear distinction: non-sales AI and general data technology contribute, at most, modestly to reducing price dispersion, whereas only sales AI has a robust and statistically significant effect on quantity dispersion—the key outcome tied to destination-level reallocation. This pattern is difficult to reconcile with a generic "digital upgrading" story and instead supports our interpretation that the primary gains operate through demand-side learning.

Using AI–Sales Hiring Intensity. In the baseline analysis, we proxy for firm-level AI–Sales adoption with a binary indicator AI^S_{ft}. As an alternative, we use a continuous measure of *AI–Sales hiring intensity*: for each firm-year, we compute the fraction of the firm's cumulative hires up to year t that are in AI–Sales positions, and use this fraction in place of the adoption dummy when constructing product-level exposure in equation (8).[38] Panel F of Table A.4 shows that the results closely mirror the baseline: products with higher exposure exhibit lower price dispersion and higher quantity dispersion.

Dropping Two Ambiguous Chinese Translations. Two keywords in our AI list—*Splunk* and *BigQuery*—do not admit clean Chinese translations. To ensure that our classification is not affected by translation noise, we exclude the Chinese translations of these terms when identifying AI-related vacancies and re-estimate the main regressions. Table A.6 shows virtually identical estimates: the coefficients are unchanged up to the third decimal place, reflecting the fact that very few postings are identified solely through these translated terms (fewer than five).

[37] We define data-related keywords as terms reflecting data infrastructure, storage, processing, and visualization tools—e.g., SQL, spreadsheets, Power BI—used primarily to manage or access information. Unlike AI-related keywords, which capture algorithmic decision-making or learning, these tools do not perform autonomous analysis or prediction.

[38] Because this fraction can be mechanically extreme when a firm has very few hires, we restrict the sample to firm-years with between 10 and 90 total hires when constructing the intensity measure.

Residualized Dispersion Measures. Finally, firms rarely make export decisions in a vacuum: even without AI, they observe and condition on common product- and market-level factors—such as global product-year shocks and destination-specific demand conditions—when forming expectations. If our dispersion measures partly reflect these shared components, it is useful to net them out so that dispersion is driven by the residual, idiosyncratic variation that remains after accounting for what is commonly observable. To do so, we construct dispersion measures using residuals after removing common components. Specifically, we residualize prices and quantities with HS6–destination and HS6–year fixed effects (and, alternatively, HS8–year fixed effects), and compute dispersion across destinations using the residuals. This procedure also helps separate a compositional channel—which firms export—from a prediction channel—how exporters allocate across markets. Table A.8 shows that our core results remain: AI exposure reduces price dispersion and increases quantity dispersion.

Classifying AI Vacancies by Occupation. To provide a systematic classification of AI-related vacancies, we implement an alternative approach based on Standard Occupational Classification (SOC) codes.[39] We match each vacancy to its SOC code and identify whether it belongs to sales, production, or other occupations. A firm is classified as adopting *sales AI* if it posts at least one AI-related vacancy in a sales occupation; similarly, *production AI* is defined using AI-related vacancies in production occupations. We then reconstruct the adoption measures and re-estimate the baseline regressions. Table A.9 reports the results, with Columns (1)–(2) focusing on sales AI and Columns (3)–(4) on production AI.

The SOC-based results closely replicate the baseline patterns. Panel A shows that sales AI adoption reduces price dispersion more strongly for products with higher initial export shares, with magnitudes very similar to the keyword-based estimates. Production AI also reduces price dispersion, but the effect is smaller, consistent with cost-efficiency improvements rather than improved market targeting. Panel B shows that sales AI adoption increases quantity dispersion, indicating more active reallocation of shipments across destinations with heterogeneous demand. In contrast, production AI has little effect on quantity dispersion: the coefficients are close to zero and statistically insignificant. Together, these findings support our interpretation that the main mechanism operates through demand-side information and destination-level allocation, rather than through production-side technology alone.

[39] We classify a vacancy as sales-related if its SOC code is 11-2011 (Advertising and Promotions Managers), 11-2021 (Marketing Managers), 13-1161 (Market Research Analysts and Marketing Specialists), or any 41-xxxx occupation (Sales and Related Occupations). We classify a vacancy as production-related if it falls under 43-5061 or any 51-xxxx occupation.

Placebo Test. To ensure that our main results are not driven by spurious patterns in the data, we conduct a placebo test by randomly reassigning the treatment variable—defined as the intensity of exposure to AI adoption—across observations. We repeat this process 500 times, constructing a new regression each time using the reshuffled treatment variable. Figure A.8 plots the resulting distribution of coefficients. Panel A shows the placebo estimates based on variation across transactions, while Panel B uses variation across destinations. In both cases, the estimated coefficients are tightly centered around zero, indicating no systematic effect from the randomly assigned treatment. By contrast, the coefficient from our baseline specification—shown as the vertical dashed line—lies far outside the placebo distribution. This stark difference confirms that the estimated effect in our main analysis is not an artifact of chance or sample structure, but it reflects a meaningful underlying relationship.

7 Conclusion

This paper develops and tests a theoretical framework in which artificial intelligence mitigates information frictions that hinder firms' ability to tailor production and pricing decisions to heterogeneous market demand. By introducing AI into a model of monopolistic competition with asymmetric information on consumer tastes, we show that firms adopting AI are better equipped to match supply with demand across destinations, resulting in more uniform prices, greater dispersion in quantities, and enhanced allocative efficiency. Empirically, we confirm these predictions using firm-level export data linked to AI-related job postings.

Our findings emphasize the value of AI-induced predictive power not just for firm performance, but also for aggregate outcomes such as price stability and consumer welfare. Importantly, the technology exerts a powerful selection effect: firms that adopt AI tend to be larger, of higher quality, and more likely to engage in exporting, reinforcing scale advantages in global markets. These results suggest that AI serves not merely as a productivity enhancer but also as an information-enabling technology that reshapes the geography of trade.

Finally, the paper anchors the emerging research agenda on AI and international trade envisioned by Goldfarb et al. (2019b). Data-rich "superstar" firms reap the lion's share of AI's gains, and we trace those gains from factoryfloor reallocations all the way to a lower aggregate cost of living. By turning a scale-and-knowledge intuition into concrete welfare improvements, the paper opens several avenues for future work—measuring AI's general-equilibrium effects on trade flows, market entry, and welfare, and giving policymakers a sharper blueprint for navigating the deepening nexus between predictive algorithms and global commerce.

References

Simona Abis and Laura Veldkamp. The changing economics of knowledge production. *The Review of Financial Studies*, 37(1):89–118, 2024.

Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Daron Acemoglu and Pascual Restrepo. Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016, 2022.

Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340, 2022.

Philippe Aghion, Simon Bunel, Xavier Jaravel, Thomas Mikaelsen, Alexandra Roulet, and Jakob Søgaard. How different uses of ai shape labor demand: evidence from france. In *AEA Papers and Proceedings*, volume 115, pages 62–67. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2025.

Ajay Agrawal, Joshua Gans, and Avi Goldfarb. *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence*. Harvard Business Press, 2022.

Ajay Agrawal, Joshua S Gans, and Avi Goldfarb. Do we want less automation? *Science*, 381(6654):155–158, 2023.

Anders Åkerman, Edwin Leuven, and Magne Mogstad. Information frictions, internet, and the relationship between distance and trade. *American Economic Journal: Applied Economics*, 14(1):133–163, 2022.

Liudmila Alekseeva, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska. The demand for ai skills in the labor market. *Labour economics*, 71:102002, 2021.

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

Alexis Antoniades, Carlos Felip Balcazar, Emmanouil Chatzikonstantinou, and Andreas Kern. Did the advent of ai offset the far-right? evidence from us' elections. *Georgetown University Working Paper*, 2025a.

Alexis Antoniades, Emmanouil Chatzikonstantinou, and Mohammed Shahmeer Ahmad. The convergence of occupations: Evidence from online job vacancies. *Georgetown University Working Paper*, 2025b.

Alexis Antoniades, Emmanouil Chatzikonstantinou, and Olesia Savka. Are ai jobs driving up the demand for stem education? yes but not for women. *Georgetown University Working Paper*, 2025c.

David Argente, Salomé Baslandze, Douglas Hanley, and Sara Moreira. Patents to products: Product innovation and firm dynamics. 2020.

Costas Arkolakis. Market penetration costs and trade dynamics. *Manuscript, Department of Economics, Yale University*, 2008.

Costas Arkolakis, Theodore Papageorgiou, and Olga A. Timoshenko. Firm learning and growth. *Review of Economic Dynamics*, 27:146–168, 2018.

Erhan Artuc, Paulo Bastos, and Bob Rijkers. Robots, tasks, and trade. *Journal of International Economics*, 145:103828, 2023.

Tania Babina. Understanding firms, households, and financing in the age of ai. Keynote lecture, CESifo-TransforM Workshop on the Economic and Societal Impacts of AI, CESifo, Munich, December 2025. URL https://www.ifo.de/en/cesifo/event/2025-12-05/cesifo-transform-workshop-economic-and-societal-impacts-ai. December 5, 2025.

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745, 2024.

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial intelligence makes firm operating performance less volatile. In *AEA Papers and Proceedings*, volume 115, pages 35–39. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2025.

Jie Bai, Maggie Chen, Jin Liu, Xiaosheng Mu, and Daniel Yi Xu. Search and information frictions on global e-commerce platforms: Evidence from aliexpress. Technical report, National Bureau of Economic Research, 2020.

Patrick Bajari, Denis Nekipelov, Stephen P. Ryan, and Miaoyu Yang. Machine learning methods for demand estimation. *The American economic review*, 105(5):481–485, 2015.

Carlos Felipe Balcazar. Globalization, unions and robots: The effects of automation on the power of labor and policymaking. *Available at SSRN 4574527*, 2023.

Paulo Bastos, Daniel A. Dias, and Olga A. Timoshenko. Learning, prices, and firm dynamics. *Canadian Journal of Economics*, 51:1257–1311, 2018.

Martin Beraja, Andrew Kao, David Y Yang, and Noam Yuchtman. Ai-tocracy. *The Quarterly Journal of Economics*, 138(3):1349–1402, 2023.

Andrew B. Bernard, J. B. Jensen, Stephen J. Redding, and Peter K. Schott. Firms in international trade. *The Journal of economic perspectives*, 21(3):105–130, 2007.

Erik Brynjolfsson. Computing productivity: are computers pulling their weight? *http://ccs. mit. edu/erik/cpg/*, 1997.

Erik Brynjolfsson and Gabriel Unger. The macroeconomics of artificial intelligence. 2023.

Jacques Bughin, Eric Hazan, Paris Sree Ramaswamy, Washington DC, Michael Chu, et al. Artificial intelligence the next digital frontier. 2017.

China Trade Post. Digital Empowerment and AI Integration—Experiencing the "Tech Vibe" at the CIFTIS. *China International Fair for Trade in Services*, 2023. URL https://www.ccpit.org/a/20230905/202309059b3q.html.

Dirk Czarnitzki, Gastón P Fernández, and Christian Rammer. Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211:188–205, 2023.

Francesco D'Acunto, Nagpurnanand Prabhala, and Alberto G Rossi. The promises and pitfalls of robo-advising. *The Review of Financial Studies*, 32(5):1983–2020, 2019.

Koen De Backer and Timothy DeStefano. Robotics and the global organisation of production. *Robotics, AI, and Humanity: Science, Ethics, and Policy*, pages 71–84, 2021.

Jan De Loecker. Do exports generate higher productivity? evidence from slovenia. *Journal of International Economics*, 73(1):69–98, 2007.

David Deming and Lisa B Kahn. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337—-S369, 2018.

David J Deming and Kadeem L Noray. Stem careers and the changing skill requirements of work. Technical report, National Bureau of Economic Research, 2019.

Michael J. Dickstein and Eduardo Morales. What do exporters know? *The Quarterly Journal of Economics*, 133:1753–1801, 2018.

Jonathan Eaton, Marcela Eslava, David Jinkins, C. J. Krizan, and James R. Tybout. A search and learning model of export dynamics, 2021.

Marius Faber. Robots and reshoring: Evidence from mexican labor markets. *Journal of International Economics*, 127:103384, 2020.

Hanming Fang, Chunmian Ge, Hanwei Huang, and Hongbin Li. Pandemics, Global Supply Chains, and Local Labor Demand: Evidence from 100 Million Posted Jobs in China. Working Paper 28072, National Bureau of Economic Research, November 2020.

Robert C Feenstra and John Romalis. International prices and endogenous quality. *The quarterly journal of economics*, 129(2):477–527, 2014.

EW Felten, Manav Raj, and Robert Seamans. The variable impact of artificial intelligence on labor: The role of complementary skills and technologies. *Available at SSRN*, 3368605, 2019.

Janos Ferencz, Javier López González, and Irene Oliván García. Artificial intelligence and international trade: Some preliminary implications. 2022.

Michael Gofman and Zhao Jin. Artificial intelligence, education, and entrepreneurship. *The Journal of Finance*, 79(1):631–667, 2024.

Avi Goldfarb, Bledi Taska, Florenta Teodoridis, et al. Could machine learning be a general-purpose technology? evidence from online job postings. *SSRN Electronic Journal*, 2019a.

Avi Goldfarb, Daniel Trefler, et al. Artificial intelligence and international trade. *The Economics of Artificial Intelligence: An Agenda*, pages 463–492, 2019b.

Jillian Grennan and Roni Michaely. Artificial intelligence and high-skilled work: Evidence from analysts. *Swiss Finance Institute Research Paper*, (20-84), 2020.

Juan Carlos Hallak and Jagadeesh Sivadasan. Product and process productivity: Implications for quality choice and conditional exporter premia. *Journal of International Economics*, 91(1):53–67, 2013.

Chuan He, Karsten Mau, and Mingzhi Xu. Trade shocks and firms hiring decisions: Evidence from vacancy postings of chinese firms in the trade war. *Labour Economics*, page 102021, 2021.

Chuan He, He Li, Hong Ma, and Mingzhi Xu. Cushioning the shock: The rise of nonstandard work amid the us-china trade war. *Working paper*, 2025.

Colin J Hottman, Stephen J Redding, and David E Weinstein. Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, 131(3):1291–1364, 2016.

Linda C. Hunter and James R. Markusen. *Per-Capita Income as a Determinant of Trade*, pages 89–109. 1988.

Leonardo Iacovone and Beata S Javorcik. Multi-product exporters: Product churning, uncertainty and export discoveries. *The Economic Journal*, 120(544):481–499, 2010.

Laurence F. Jackson. Hierarchic demand and the engel curve for variety. *The review of economics and statistics*, 66(1):8–15, 1984.

Robert Jensen. The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The Quarterly journal of economics*, 122(3): 879–924, 2007.

Lisa B. Kahn and Brad Hershbein. Do recessions accelerate routine-biased technological change? *American Economic Review*, 108(7):1737–1772, 2018.

Chan Kim and Daisoon Kim. Trade search frictions: Public information, exports, and concentration. *Exports, and Concentration*, 2024.

Astrid Krenz, Klaus Prettner, and Holger Strulik. Robots, reshoring, and the lot of low-skilled workers. *European Economic Review*, 136:103744, 2021.

Maurice Kugler and Eric Verhoogen. Prices, plant size, and product quality. *The Review of Economic Studies*, 79(1):307–339, 2012.

Peter Kuhn and Kailing Shen. Gender discrimination in job ads: Evidence from china. *The Quarterly Journal of Economics*, 128(1):287–336, 2013.

Lucia Mỳtna Kureková, Miroslav Beblavỳ, and Anna Thum-Thysen. Using online vacancies and web surveys to analyse the labour market: A methodological inquiry. *IZA Journal of Labor Economics*, 4:1–20, 2015.

Magnus Lodefalk, Aili Tang, and Miaojie Yu. Stayin'alive: Export credit guarantees and export survival. Technical report, Working Paper, 2022.

Luca Macedoni. Asymmetric information, quality, and regulations. *Review of International Economics*, 30(4):1180–1198, 2022.

Beatrice Magistro, Sophie Borwein, R. Michael Alvarez, Bart Bonikowski, and Peter John Loewen. Attitudes toward artificial intelligence (ai) and globalization: Common micro-foundations and political implications. *American Journal of Political Science*, 2025.

Kalina Manova and Zhiwei Zhang. Export prices across firms and destinations. *The Quarterly Journal of Economics*, 127(1):379–436, 2012.

Tshilidzi Marwala, Eleonore Fournier-Tombs, and Serge Stinckwich. The use of synthetic data to train ai models: Opportunities and risks for sustainable development. *arXiv preprint arXiv:2309.00652*, 2023.

Kristina Mcelheran, Mu Jeung Yang, Erik Brynjolfsson, and Zachary Kroff. The rise of industrial ai in america. 2025.

Jakob R. Munch and Daniel X. Nguyen. Decomposing firm-level sales variation. *Journal of Economic Behavior Organization*, 106:317–334, 2014.

Ozcan Ozturk. The impact of ai on international trade: Opportunities and challenges. *Economies*, 12(11):298, 2024.

Sheetal Pansare. Predictive Analytics in Global Trade: Forecasting Market Trends with AI. *Global Trade*, 2024. URL https://www.globaltrademag.com/predictive-analytics-in-global-trade-forecasting-market-trends-with-ai/.

Manav Raj and Robert Seamans. Artificial intelligence, labor, productivity, and the need for firm-level data. In *The economics of artificial intelligence: An agenda*, pages 553–565. University of Chicago Press, 2018.

Kelly Raskovich, Bill Briggs, and Mike Bechtel. Predictive Analytics in Global Trade: Forecasting Market Trends with AI. *Deloitte Insights*, 2025. URL https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends.html.

Erick Sager and Olga A. Timoshenko. Demand uncertainty, selection, and trade. *FEDS Working Paper No. 2024-42*, 2024.

Ina Simonovska. Income differences and prices of tradables: Insights from an online retailer. *The Review of Economic Studies*, 82(4 (293)):1612–1656, 2015.

Alex Singla, A Sukharevsky, L Yee, M Chui, and B Hall. The state of AI: How organizations are rewiring to capture value. *McKinsey & Company*, 12, 2025. URL https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai.

Katherine Stapleton and Michael Webb. Automation, trade and multinational activity: Micro evidence from spain. *Available at SSRN 3681143*, 2020.

Miroslav Štefánik. Internet job search data as a possible source of information on skills demand (with results for slovak university graduates). *Building on skills forecasts—Comparing methods and applications*, 246, 2012.

Stephanie Steinmetz, Damian Raess, Kea Tijdens, and Pablo de Pedraza. Measuring wages worldwide: exploring the potentials and constraints of volunteer web surveys. *Advancing research methods with new technologies*, pages 100–119, 2013.

Claudia Steinwender. Real effects of information frictions: When the states and the kingdom became united. *American Economic Review*, 108(3):657–696, 2018.

Linnet Taylor, Ralph Schroeder, and Eric Meyer. Emerging practices and perspectives on big data analysis in economics: Bigger and better or more of the same? *Big Data & Society*, 1(2):2053951714536877, 2014.

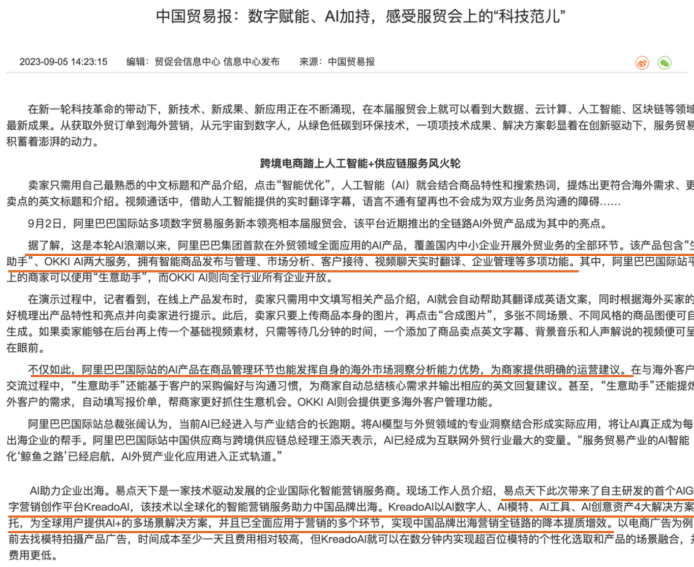
Johannes Van Biesebroeck. Exporting raises productivity in sub-saharan african manufacturing firms. *Journal of International Economics*, 67(2):373–391, 2005.

Johannes Van Biesebroeck and Elena Zaurino. Unilateral trade preferences and export growth in developing countries. *The World Bank Economic Review*, page lhaf015, 2025.

Joachim Wagner. Temporary exports and characteristics of destination countries: first evidence from german transaction data. *Economics*, 12(1):20180054, 2018.

Appendix

Figure A.1: Reports on AI Marketing and Intelligent Exporting



中国贸易报：数字赋能、AI加持，感受服贸会上的“科技范儿”

2023-09-05 14:23:15　编辑：贸促会信息中心 信息中心发布　来源：中国贸易报

在新一轮科技革命的带动下，新技术、新成果、新应用正在不断涌现，在本届服贸会上就可以看到大数据、云计算、人工智能、区块链等领域的最新成果。从获取外贸订单到海外营销，从元宇宙到数字人，从绿色低碳到环保技术，一项项技术成果、解决方案彰显着在创新驱动下，服务贸易正积蓄着澎湃的动力。

跨境电商踏上人工智能+供应链服务风火轮

卖家只需用自己最熟悉的中文标题和产品介绍，点击“智能优化”，人工智能（AI）就会结合商品特性和搜索热词，提炼出更符合海外需求、更有卖点的英文标题和介绍。视频通话中，借助人工智能提供的实时翻译字幕，语言不通有望再也不会成为双方业务员沟通的障碍……

9月2日，阿里巴巴国际站多项数字贸易服务新本领亮相本届服贸会，该平台近期推出的全链路AI外贸产品成为其中的亮点。

握了解，这是本轮AI浪潮以来，阿里巴巴集团首款在外贸领域全面应用的AI产品，覆盖国内中小企业开展外贸业务的全部环节。该产品包含“生意助手”、OKKI AI两大服务，拥有智能商品发布与管理、市场分析、客户接待、视频聊天实时翻译、企业管理等多项功能。其中，阿里巴巴国际站平台上的商家可以使用“生意助手”，而OKKI AI则向全行业所有企业开放。

在演示过程中，记者看到，在线上产品发布时，卖家只需用中文填写相关产品介绍，AI就会自动帮助其翻译成英语文案，同时根据海外买家的喜好梳理出产品特性和亮点并向卖家进行提示。此后，卖家只要上传商品本身的图片，再点击“合成图片”，多张不同场景、不同风格的商品图便可自动生成。如果卖家能够在后台再上传一个基础视频素材，只需等待几分钟的时间，一个添加了商品卖点英文字幕、背景音乐和人声解说的视频便可呈现在眼前。

不仅如此，阿里巴巴国际站的AI产品在商品管理环节也能发挥自身的海外市场洞察分析能力优势，为商家提供明确的运营建议。在与海外客户的交流过程中，“生意助手”还能基于客户的采购偏好与沟通习惯，为商家自动总结核心需求并输出相应的英文回复建议。甚至，“生意助手”还能提炼海外客户的需求，自动填写报价单，帮商家更好抓住生意机会。OKKI AI则会提供更多海外客户管理功能。

阿里巴巴国际站总裁张阎认为，当前AI已经进入与产业结合的长跑期。将AI模型与外贸领域的专业洞察结合形成实际应用，将让AI真正成为每个出海企业的帮手。阿里巴巴国际站中国供应商与跨境供应链总经理王添天表示，AI已经成为互联网外贸行业最大的变量。“服务贸易产业的AI智能化‘鲸鱼之路’已经启航，AI外贸产业化应用进入正式轨道。”

AI助力企业出海。易点天下是一家技术驱动发展的企业国际化智能营销服务商。现场工作人员介绍，易点天下此次带来了自主研发的首个AIGC数字营销创作平台KreadoAI，该技术以全球化的智能营销服务助力中国品牌出海。KreadoAI以AI数字人、AI模特、AI工具、AI创意资产4大解决方案为依托，为全球用户提供AI+的多场景解决方案，并且已全面应用于营销的多个环节，实现中国品牌出海营销全链路的降本提质增效。以电商广告为例，以前去找模特拍摄产品广告，时间成本至少一天且费用相对较高，但KreadoAI就可以在数分钟内实现超百位模特的个性化选取和产品的场景融合，并且费用更低。

China Trade News: Digital Empowerment and AI Integration—Experiencing the “Tech Vibe” at the CIFTIS

According to reports, this is *Alibaba*’s first AI product fully applied to the foreign trade sector since the latest wave of AI advancements. The product covers all aspects of foreign trade operations for domestic SMEs. It includes two main services, “Business Assistant” and “*OKKI AI*,” offering a range of features such as intelligent product listing and management, market analysis, customer engagement, real-time video chat translation, and enterprise management.

Moreover, *Alibaba* International’s AI product also leverages its market intelligence and analytical capabilities in product management, providing merchants with clear and actionable operational insights for overseas markets.

Yi Dian Tian Xia has introduced *KreadoAI*, its first self-developed AIGC digital marketing creation platform, designed to support Chinese brands in expanding globally through AI-powered marketing services. Built on four core solutions—AI digital humans, AI models, AI tools, and AI creative assets—*KreadoAI* provides multi-scenario AI+ solutions for global users. The technology has been fully integrated into various marketing stages, enhancing efficiency, reducing costs, and improving quality across the entire marketing chain for Chinese brands entering international markets.

Note: Authors’ translation of China Trade Post (2023), published by *China Council For the Promotion of International Trade*.

Table A.1: Keywords for Identifying AI and Data Vacancy

Panel A: AI-Related Keywords			*Panel B: Data-Related Keywords*		
Keywords Index	Chinese	English	Keywords Index	Chinese	English
1	人工智能	artificial intelligence	1	数据分析	data analysis
	机器智能			数据透视	
	虚拟智能			分析数据	data analytics
	模拟智能			透视数据	
2	自动化	automation	2	数据建模	data modeling
	自动技术			数据模型	
3	机器学习	machine learning	3	数据工程	data engineering
4	深度学习	deep learning	4	数据可视	data visualization
5	人工超级智能	artificial superintelligence	5	数据提取	extract data
	人工超智能			提取数据	
	虚拟超级智能			信息提取	
	虚拟超智能			提取信息	
	模拟超级智能			数据获取	
	模拟超智能			获取数据	
6	张量流	tensorflow	6	数据科学	data science
7	自然语言处理	natural language processing	7	商务智能	power bi
		nlp		商业智能	powerbi
8	决策树	decision tree	8	电子表格	spreadsheet
9	贝叶斯	bayes	9	数据仓库	etl
10	神经网络	neural networks	10	数据管理	data management
		cnn		管理数据	
11	监督算法	supervised algorithm	11	数据挖掘	data mining
	监督学习	supervised learning		信息挖掘	
12	分布式计算	mapreduce		挖掘数据	
	并行式计算			挖掘信息	
	并行计算		12	预测分析	predictive analysis
13	机器数据引擎	splunk		商业分析	
	机器数据的引擎			预测技术	
14	聚类分析	cluster analysis	13	预测模型	predictive model
15	支持向量	support vector		预测建模	
	向量机	svm		预测模式	
16	阿帕奇	apache	14	数智化	tebleau
	阿帕切			报表分析	
17		hadoop		表格分析	
18	计算引擎	spark	15	数据转换	convert data
19		kubernetes		转换数据	
		k8s		数据变换	
20		torch		变换数据	
21	大查询	bigquery	16	数据下载	download data
	云端存储			下载数据	
	云端储存			加载数据	
	云端资料分析			数据加载	
22	语言模型	language mode	17	数据库	sql
	语言模组			结构化查询语言	
	语言模式	llm		结构查询语言	
23	计算机视觉	computer vision		数据查询语言	
	机械视觉		18	分布式存储	mongodb
	机器视觉			分布式文件存储	
24	扩散模型	diffusion model		分布式储存	
	扩散模式			分布式文件储存	
	传播模型				

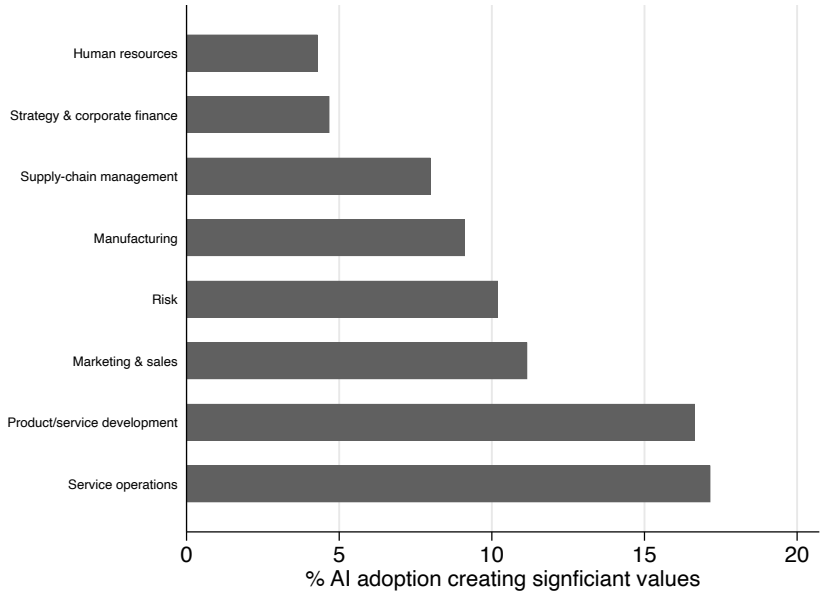
Note: This table lists the keywords used to classify job postings as AI-related or data-technology-related. In the implementation, we search both the English keywords and their corresponding Chinese translations. Panel A (AI-related) includes terms associated with algorithmic prediction, machine learning, and automated optimization—tools designed to replace or augment human judgment by learning patterns from data, including early-generation AI methods. Panel B (data-related) includes terms referring to data infrastructure and analytics tooling for storing, processing, managing, or visualizing data. These keywords capture investments in data handling and analytical capacity rather than AI-driven automation or predictive intelligence.

Table A.2: Industry Categories Used in the Study

Index	Industry
1	Transportation, storage and postal services
2	Accommodation and catering
3	Information transmission, software and information technology services
4	Public administration, social security and social organizations
5	Agriculture, forestry, animal husbandry and fishery
6	Manufacturing
7	Health and social work
8	International organization
9	Residential services, repair and other services
10	Construction
11	Real estate
12	Wholesale and retail
13	Education
14	Culture, sports and entertainment
15	Water conservancy, environment and public facilities management
16	Electricity, heat, gas and water production and supply
17	Scientific research and technology services
18	Leasing and business services
19	Mining
20	Finance

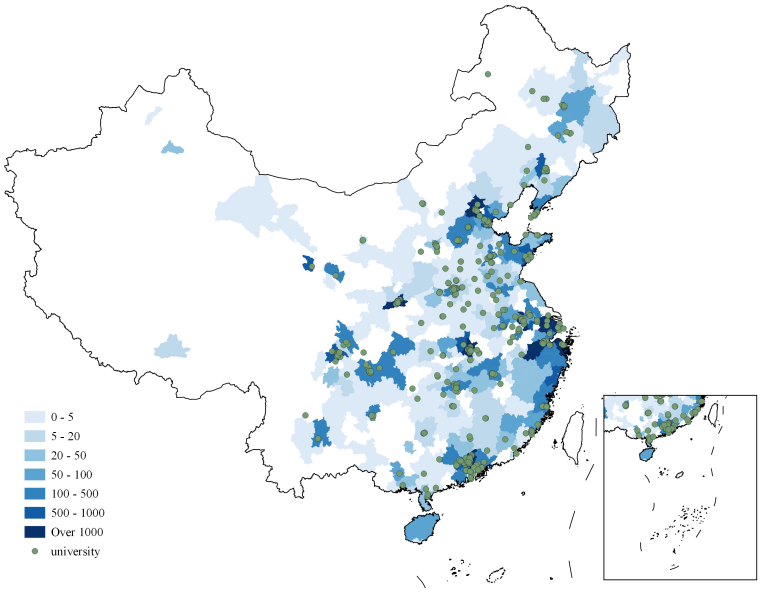
Note: The listed industries correspond to the 20 top-level sections in the Industrial Classification for National Economic Activities (GB/T 4754-2017) published by the Standardization Administration of China. Lower-tier codes—division (2-digit), group(3-digit), and class (4-digit)—are aggregated to the section level in this study.

Figure A.2: Business Functions Where Firms Report Value from AI Adoption



Notes: Authors' calculation based on McKinsey Global Institute's report on AI adoptions, surveying more than three thousand companies across nations in 2017. Each bar denotes the share of companies that self-reported significant value gain from adopting the corresponding AI function. The details of the survey can be found in Bughin et al. (2017).

Figure A.3: Spatial Distribution of AI-Related University Enrollment Quotas and Firm Hiring Intensity



Note: This figure illustrates the geographic distribution of AI-related university enrollment quotas alongside firm-level AI hiring demand. Green dots denote universities with substantial AI-related enrollment (exceeding 300 students). The shading represents the density of AI-related job vacancies by prefecture, reflecting regional hiring intensity. The spatial overlap between university talent supply (green dots) and firm demand (shaded regions) suggests a strong correlation between the localization of AI education and labor market needs.

A Data Representativeness

We provide a comprehensive discussion of the representativeness of our data by comparing it with other representative data, pointing out potential limitations, and summarizing the nature of our data and its implications.

Comparison with administrative data. We compare key job characteristics—wages, required work experience, and the share of positions requiring a college degree—to national averages derived from administrative records. Table A.3 presents these comparisons, using data from the 2014–2016 OJV sample alongside national statistics from the China Statistical Yearbook and Population Census. Consistent with findings in the literature on OJV data (Kuhn and Shen, 2013; Deming and Kahn, 2018; Kahn and Hershbein, 2018; Acemoglu et al., 2022), our sample skews toward high-skilled, higher-paid jobs, with offered wages approximately 30% higher and the share of position requiring a college degree doubled that of the administrative statistics. This aligns with the nature of online job postings, which tend to capture demand for skilled labor in high-tech and emerging sectors more effectively than traditional administrative records. In particular, He et al. (2025) document the prevalence of gig, skilled work in Chinese OJV data, driven by the development of internet platforms.

Table A.3: Job characteristics: our Samples vs. administrative data

Variables	(1) All OJV	(2) Matched 14-16	(3) Matched 14-15	(4) Baseline	(5) Administrative
Offered Wage (10,000 RMB per year)	10.18	10.21	10.22	10.30	6.75
	(5.591)	(5.519)	(5.526)	(5.560)	
Working Experience (year)	2.121	2.385	2.348	2.355	–
	(2.192)	(2.236)	(2.229)	(2.236)	
Share of college degree (%)	70.6%	73.2%	72.8%	72.7%	32.23%
	(0.455)	(0.443)	(0.445)	(0.446)	

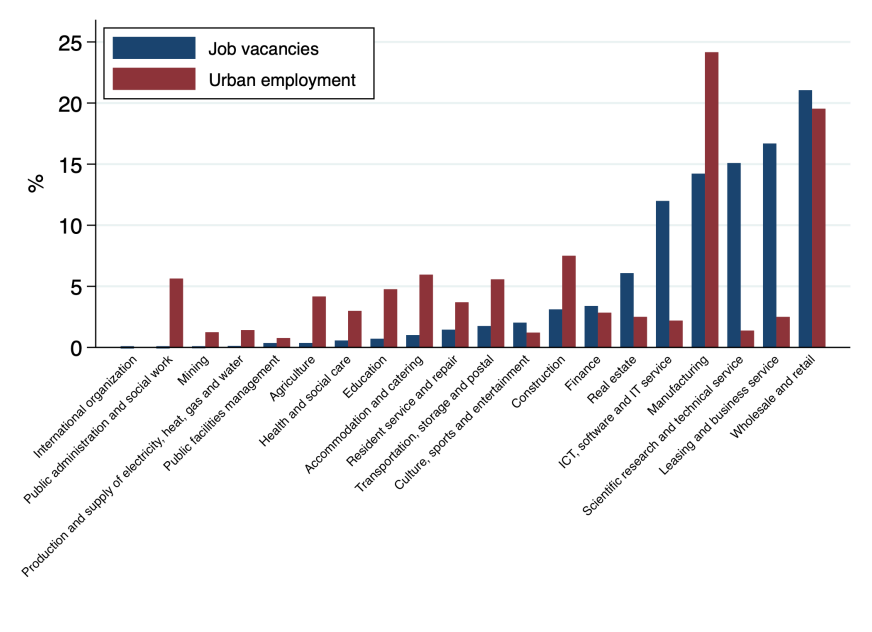
Note: Columns (1)–(4) report statistics for different subsamples: (1) the full OJV dataset, (2) a matched sample including only firms engaged in trade between 2014 and 2016, (3) a subset with additional trade quantity and price information (2014–2015), and (4) the sample used in our baseline model. Column (5) presents administrative statistics for the broader labor market in 2016. The Administrative data in column (5) are sourced from the China Statistical Yearbook (2017), which reports the average wage level of 2016, and Population Census (2015), which documents the urban employed population with some college education.

Despite this skewness, these data characteristics are advantageous for our study, as AI skill adoption is more prevalent in high-wage, high-tech sectors and among firms engaged in international trade. The overrepresentation of such industries aligns with the nature of AI diffusion, which tends to occur first in technologically advanced and knowledge-intensive

sectors. Moreover, exporting firms—often larger, more productive, and globally competitive—are more likely to invest in AI-driven innovations to enhance efficiency and maintain a competitive edge in international markets.

Figure A.4 compares the industrial distribution of job vacancies from 2014–2016 with the realized urban employment from administrative records. Our dataset provides a strong representation of the three largest industries—manufacturing, wholesale and retail, and construction—while exhibiting a bias toward high-tech and business-related sectors. Consistent with prior research on OJV data, our sample overrepresents industries such as ICT, finance, scientific research, and leasing and business services, where demand for skilled labor and AI-related technologies is more pronounced. In contrast, industries such as public administration, mining, health and social care, and education are underrepresented, reflecting the limited reliance on OJV platforms for recruitment in these sectors. This sectoral composition aligns well with our research focus. AI adoption is more likely to emerge in technology-intensive, service-oriented, and internationally competitive industries, where firms have both the incentive and capacity to integrate AI-driven innovations.

Figure A.4: Industry distribution of OJV

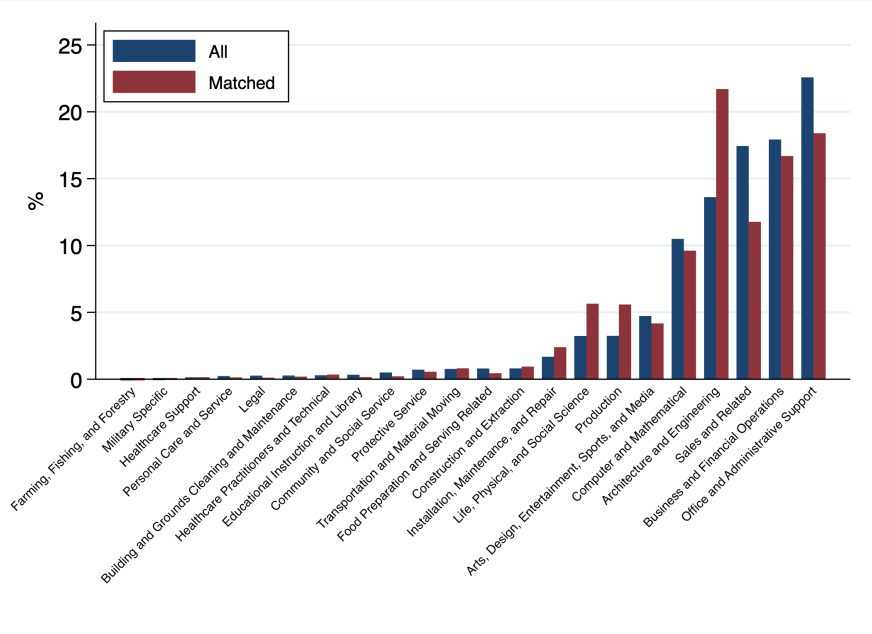


Note: Authors' calculation based on sampled job vacancies during 2014-2016. Red bars indicate the shares of urban employment by industry documented in the Population Census (2015).

We match our data to the O*NET occupation classification to check the occupational composition of OJV data. Figure A.5 presents the occupational distribution of all job vacancies in our dataset, compared to a matched subset of firms engaged in trade. The distribution

highlights a strong presence of business-related and high-skilled occupations, with a notable concentration in office and administrative support, business and financial operations, sales, and computer and mathematical roles. These occupations align well with our research focus, as AI adoption is more likely to reshape tasks in knowledge-intensive and data-driven professions.

Figure A.5: Occupational distribution of OJV



Note: Authors' calculation based on sampled job vacancies during 2014-2016. Blue bars indicate the whole OJV sample and red bars indicate the subset of firms engaged in trade. The occupation classification is based on O*NET's *job family* categories. The Management category is excluded to improve comparability across non-managerial occupations, given its broad and heterogeneous definition in Chinese job titles.

Notably, computer and mathematical occupations are well represented, which is critical for studying AI skill adoption, as these fields are at the forefront of AI development and implementation. Additionally, sales, financial operations, and administrative support roles play an essential role in firms' export behaviors, particularly in managing trade logistics, market expansion, and customer relationship management. The prevalence of these occupations in our dataset enhances its suitability for examining AI's impact on firms operating in international markets.

From an export perspective, the dataset remains highly relevant, as many of the represented occupations are critical to firms' global competitiveness. AI-related transformations in administrative support, finance, and sales functions could directly influence firms' ability to manage trade relationships, optimize supply chains, and scale operations in international

markets. The high representation of computer and mathematical occupations further enhances the suitability of our data, as these roles are at the frontier of AI adoption.

Limitation. While the OJV data provide broad coverage and granular insights into job vacancies, we acknowledge certain limitations regarding their representativeness of the overall labor market, which warrant some caution in interpreting our findings. First, consistent with prior studies, online job postings tend to overrepresent high-skilled, high-wage jobs, disproportionately favoring educated workers and industries at the technological frontier, such as ICT, finance, and business services. Traditional sectors, including public administration, education, health care, and low-tech manufacturing, are underrepresented, limiting our ability to generalize findings to these segments of the labor market. Second, occupational representation is also skewed, with a concentration in business-related, administrative, and technical roles, while sectors that traditionally rely on lower-skilled employment—such as construction, food services, and manual labor—are less visible in the dataset. Third, strategic hiring practices may introduce selection bias, as firms may not post all job openings online. If firms selectively advertise certain job types online, our results may reflect differential hiring behaviors rather than underlying changes in labor demand.

Summary. Online data and web-based surveys share a natural limitation that the offline segment of information often falls out of the sample (Štefánik, 2012; Steinmetz et al., 2013; Taylor et al., 2014). Specifically, the unobservable population and structure of job vacancies make it challenging to address representativeness issues of OJV data (Kureková et al., 2015). As highlighted in the relevant literature and shown in this section, OJV data is likely to represent a segment of the labor market with higher skill levels and in emerging sectors.

Despite these caveats, OJV data remains one of the best available sources for studying the intersection of AI adoption, skill demand, and firm-level trade behavior in China. AI-driven technological change and export-oriented firms are heavily concentrated in industries where online recruitment is a primary hiring channel, making this dataset well-suited for examining how firms adjust their workforce in response to AI adoption and global trade pressures. While challenges related to representativeness persist, we mitigate these concerns through robustness checks, including excluding specific regions and using alternative measures, to validate our results. By leveraging OJV data, which disproportionately captures job postings from firms at the frontier of technological adoption, our study can more effectively analyze the demand for AI-related skills and the impact of AI technology on firms' employment structures and export behaviors.

B Model Appendix

B.1 The Autarchy Case without AI

B.1.1 Firms' Pricing Decision

After production, each firm chooses its price to maximize profits:

$$\max_{p} \; p \min \left\{ q, \frac{z^{\sigma-1} s^{\sigma-1} p^{-\sigma}}{P^{1-\sigma}} \right\}$$

There are two possible cases. If $p \geq (zsP)^{\frac{\sigma-1}{\sigma}} q^{-\frac{1}{\sigma}}$, then demand is less than capacity. In this condition, lowering the price increases sales and thus raises profits. If $p \leq (zsP)^{\frac{\sigma-1}{\sigma}} q^{-\frac{1}{\sigma}}$, then production capacity is binding. In this case, raising the price while selling the same fixed quantity increases profits. Putting the two cases together, the firm's optimal price is given by the cutoff value: $p = (zsP)^{\frac{\sigma-1}{\sigma}} q^{-\frac{1}{\sigma}}$.

B.1.2 Firms' Quantity Decision

Having solved for the optimal price in the third sub-period, we now move backward to the firm's quantity choice. Substituting the optimal price, along with the variable cost, into the firm's profit function yields the following problem:

$$\max_{q} \; \mathbb{E}_z((zsP)^{\frac{\sigma-1}{\sigma}} q^{-\frac{1}{\sigma}} - s^{\alpha})q - f.$$

Taking the first-order condition with respect to q and solving gives the optimal output as a function of s:

$$q(s) = \left[\frac{\sigma - 1}{\sigma} s^{\frac{\sigma-1}{\sigma} - \alpha} P^{\frac{\sigma-1}{\sigma}} \mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}}) \right]^{\sigma}.$$

B.1.3 Firms' Entry Decision

Along the equilibrium path, firms set their price as

$$p(s, z) = \frac{\sigma}{\sigma - 1} \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} s^{\alpha}.$$

The resulting *ex post* profit (net of the entry cost) is given by

$$\pi(s, z) = \left(\frac{\sigma}{\sigma - 1} \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} - 1 \right) \left[\frac{\sigma - 1}{\sigma} \mathbb{E}_z z^{\frac{\sigma-1}{\sigma}} P^{\frac{\sigma-1}{\sigma}} \right]^{\sigma} s^{(\sigma-1)(1-\alpha)} - f.$$

Taking expectations over z, the expected profit for a firm with quality s is

$$\mathbb{E}_z(\pi(s,z)) = \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{\sigma} \left[(\sigma-1)s^{1-\alpha}P\right]^{\sigma-1} - f.$$

The cutoff s_p is defined by the zero-profit condition, so that a firm produces if and only if its quality index exceeds s_p:

$$s_p = \frac{f^{\frac{1}{(\sigma-1)(1-\alpha)}}}{((\sigma-1)P)^{\frac{1}{1-\alpha}}} \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{(\sigma-1)(1-\alpha)}}.$$

Expected profits can then be expressed in terms of this cutoff:

$$\mathbb{E}_z(\pi(s,z)) = \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f.$$

Finally, the free-entry condition requires that the expected gains from entry equal the entry cost f_e:

$$\int_{s_p}^{+\infty} (\mathbb{E}_z(\pi(s,z)))dF(s) = f_e,$$

or equivalently,

$$\int_{s_p}^{+\infty} \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f dF(s) = f_e.$$

This condition uniquely determines the cutoff s_p, which characterizes the minimum quality level required for firms to profitably produce.

B.1.4 Equilibrium

In equilibrium, the labor market clearing condition requires that the total demand for labor equals the available supply. This condition can be written as

$$N_e \left[\int_{s_p}^{+\infty} (q(s)s^{\alpha} + f)dF(s) + f_e\right] = 1$$

where N_e denotes the number of entrants. Using the relationship $q(s)s^{\alpha} = (\sigma-1)(\mathbb{E}_z(\pi(s,z)) + f)$, we can express labor demand as

$$N_e \int_{s_p}^{+\infty} \left[(\sigma-1)\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} + 1\right] f dF(s) + N_e f_e = 1.$$

Combining this with the free-entry condition, the equilibrium number of entrants satisfies

$$N_e \int_{s_p}^{+\infty} \left[\sigma \left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} \right] f dF(s) = 1.$$

Hence, the equilibrium mass of entrants is given by

$$N_e = \frac{1}{\sigma f \int_{s_p}^{+\infty} \left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} dF(s)}.$$

Finally, using the zero-profit condition at the cutoff s_p, the equilibrium price index can be written as

$$P = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{1-\alpha}} \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma} \right]^{-\frac{\sigma}{\sigma-1}}.$$

Together, these conditions characterize the equilibrium of the economy.

Definition 1 (Equilibrium without AI) *Given* $\{F(s), G(z), \sigma, f, f_e\}$*, an equilibrium consists of* $\{q(s), p(s,z), s_p, N_e, P\}$ *such that*

(1) customers maximize their utility;

(2) firms maximize their (expected) profit, which implies $p(s,z) = \frac{\sigma}{\sigma-1} \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} s^{\alpha}$ *and* $q(s) = \left[\frac{\sigma-1}{\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} P^{\frac{\sigma-1}{\sigma}} \mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}}) \right]^{\sigma}$*;*

(3) zero expected profit (excluding entry cost) at the cutoff

$$\mathbb{E}_z(\pi(s_p, z)) = 0;$$

(4) free entry condition is satisfied

$$\int_{s_p}^{+\infty} (\mathbb{E}_z(\pi(s,z))) dF(s) = f_e;$$

(5) goods market clears;

(6) labor market clears

$$N_e \int_{s_p}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f dF(s) + N_e f_e = 1;$$

(7) price index is given by

$$P = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{1-\alpha}}\left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.$$

B.1.5 Proof of the Existence and Uniqueness of Equilibrium

The solution to consumers' utility maximization problem is given by the demand function. Firms' profit maximization problem is solved jointly by their pricing decision $p(s,z)$, quantity decision $q(s)$, and production decision, i.e., the cutoff strategy s_p with zero expected profit at the cutoff. Note that the price set by the firm exactly clears the goods market.

We then first show that there exists a unique s_p that satisfies the free entry condition. Note that

$$h(s_p) \equiv \int_{s_p}^{+\infty}\left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f dF(s)$$

is a continuous function of s_p on the domain of $(0,+\infty)$ and is strictly decreasing in s_p. What's more, $\lim_{s_p\to 0} h(s_p) = +\infty$ and $\lim_{s_p\to\infty} h(s_p) = 0$. We can then conclude that there exists a unique $s_p > 0$ that satisfies the free entry condition $\int_{s_p}^{+\infty}\left(\mathbb{E}_z(\pi(s,z))\right)dF(s) = f_e$. The price index implies that zero expected profit condition at the cutoff satisfies. Finally, the number of entrants is given by the labor market clear condition

$$N_e = \frac{1}{\sigma f\int_{s_p}^{+\infty}\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} dF(s)},$$

as we have solved for s_p.

To complete the proof for the existence of the equilibrium, we need to verify that the price index given by the zero expected profit condition is exactly what consumers face from

the demand side. To see this, we derive

$$
\begin{aligned}
P &= \left[\int_{\omega\in\Omega}\left(\frac{p(\omega)}{z(\omega)s(\omega)}\right)^{1-\sigma}d\omega\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e\int_{s_p}^{+\infty}\mathbb{E}_z\left(\frac{p(s,z)}{zs}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e\int_{s_p}^{+\infty}\mathbb{E}_z\left(\frac{\sigma}{\sigma-1}\frac{z^{-\frac{1}{\sigma}}}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}s^{\alpha-1}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\frac{1}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}\left[N_e\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\int_{s_p}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\frac{1}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}\left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)s_p^{(\sigma-1)(1-\alpha)}}{\sigma f}\right]^{\frac{1}{1-\sigma}} \\
&= \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{1-\alpha}}\left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.
\end{aligned}
$$

This is exactly what we define as the price index. We complete the proof of the existence of the equilibrium. What's more, as the solution for s_p is unique, the uniqueness of the equilibrium thus follows. Now we finish the proof for the existence and uniqueness of the equilibrium under asymmetric information.

B.2 The Autarchy Case with AI

B.2.1 Firms' AI Deployment Decision

Without AI, the expected profit of a firm before knowing consumers' taste is

$$
\mathbb{E}_z(\pi^{AI}(s,z|n)) = \frac{1}{\sigma-1}\left[\frac{\sigma-1}{\sigma}\mathbb{E}_z\Big(z^{\frac{\sigma-1}{\sigma}}\Big)P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma}s^{(\sigma-1)(1-\alpha)} - f.
$$

On the other hand, the expected profit after AI use is

$$
\mathbb{E}_z(\pi^{AI}(s,z|a)) = \frac{\mathbb{E}_z(z^{\sigma-1})}{\sigma^{\sigma}}\Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1} - f - f_{AI}.
$$

Therefore, the firm would choose to adopt AI if and only if

$$\mathbb{E}_z(\pi^{AI}(s,z|a)) \geq \mathbb{E}_z(\pi^{AI}(s,z|n))$$

i.e.,

$$s \geq \left[\frac{\sigma^\sigma f_{AI}}{[(\sigma-1)P]^{\sigma-1}\left(\mathbb{E}_z(z^{\sigma-1}) - \mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)^\sigma\right)}\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}.$$

Note that from Jensen's inequality, we know $\mathbb{E}_z(z^{\sigma-1}) - \mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)^\sigma \geq 0$ since $\sigma > 1$. Also, the equality will not hold as the support of $G(z)$ is on the interval $[z_l, z_h]$. Thus, the right hand side of the above equation is larger than zero. We conclude that firms would adopt a cutoff strategy for AI use with $s_a^{AI} = \left[\frac{\sigma^\sigma f_{AI}}{[(\sigma-1)P]^{\sigma-1}\left(\mathbb{E}_z(z^{\sigma-1}) - \mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)^\sigma\right)}\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}$.

B.2.2 Threshold Comparison

Comparing s_a^{AI} with $s_p^{AI} = \frac{f^{\frac{1}{(\sigma-1)(1-\alpha)}}}{\left((\sigma-1)P\right)^{\frac{1}{1-\alpha}}}\left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{-\frac{\sigma}{(\sigma-1)(1-\alpha)}}$, we have

$$s_a^{AI} = s_p^{AI}\left(\frac{f_{AI}\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^\sigma}{f\left(\mathbb{E}_z(z^{\sigma-1}) - \left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^\sigma\right)}\right)^{\frac{1}{(\sigma-1)(1-\alpha)}} > s_p^{AI}$$

by Assumption 2, i.e., the cutoff for using AI is higher than that for production.

B.2.3 Equilibrium

From the zero expected profit condition at the production cutoff and indifference condition at the AI cutoff, the expected profit for firms with quality $s \in (s_p^{AI}, s_a^{AI})$ can be written as

$$\mathbb{E}_z(\pi^{AI}(s,z|n)) = \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f,$$

and that for firms with quality $s \in [s_a^{AI}, +\infty)$ can be written as

$$\mathbb{E}_z(\pi^{AI}(s,z|a)) = \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f + \left[\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f_{AI}.$$

From free entry condition, we have

$$\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z \left(\pi^{AI}(s,z)\right) dF(s) = f_e,$$

i.e.,

$$\int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f dF(s) + \int_{s_a^{AI}}^{+\infty} \left[\left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI} dF(s) = f_e.$$

With the relationship between s_p^{AI} and s_a^{AI}, we can solve for s_p^{AI} and s_a^{AI} from the equations. Note that s_a^{AI} increases with s_p^{AI}.

We now define the equilibrium with access to AI technology.

Definition 2 (Equilibrium with AI) *Given $\{F(s), G(z), \sigma, f, f_e\}$, an equilibrium consists of $\{q^{AI}(s,z), p^{AI}(s,z), s_p^{AI}, s_a^{AI}, N_e^{AI}, P\}$ such that*

(1) customers maximize their utility;

(2) firms maximize their (expected) profit, which implies $p^{AI}(s,z) = \frac{\sigma}{\sigma-1} s^{\alpha}$ and $q^{AI}(s,z) = \left[\frac{\sigma-1}{\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} z^{\sigma-1}$ for firms with $s \in [s_a^{AI}, +\infty)$, and $p^{AI}(s,z) = \frac{\sigma}{\sigma-1} \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} s^{\alpha}$ and $q^{AI}(s) = \left[\frac{\sigma-1}{\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} P^{\frac{\sigma-1}{\sigma}} \mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}$ for firms with $s \in [s_p^{AI}, s_a^{AI})$;

(3) zero expected profit (excluding entry cost) at the cutoff

$$\mathbb{E}_z(\pi^{AI}(s_p^{AI}, z|n)) = 0;$$

(4) indifference about the expected profit between using AI and not using at the AI cutoff

$$\mathbb{E}_z(\pi^{AI}(s_a^{AI}, z|n)) = \mathbb{E}_z(\pi^{AI}(s_a^{AI}, z|a));$$

(5) free entry condition is satisfied

$$\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z(\pi^{AI}(s,z)) dF(s) = f_e;$$

(6) goods market clears;

(7) labor market clears

$$N_e^{AI} \left\{ \int_{s_p^{AI}}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f dF(s) + \int_{s_a^{AI}}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f_{AI} dF(s) + f_e \right\} = 1;$$

(8) price index is given by

$$P^{AI} = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{AI\ 1-\alpha}} \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.$$

B.2.4 Proof of the Existence and Uniqueness of Equilibrium

The solution to consumers' utility maximization problem is given by the demand function. Firms' profit maximization problem is solved jointly by their pricing decision $p^{AI}(s,z)$, quantity decision $q^{AI}(s,z)$, and production decision and AI adoption decision, i.e., the cutoff strategy s_p^{AI} with zero expected profit at the cutoff and s_a^{AI} with indifferent expected profit at the cutoff. Note that the price set by the firm exactly clears the goods market.

We then first show that there exists a unique set of s_p^{AI} and s_a^{AI} that satisfies the free entry condition. By defining $C \equiv \left(\frac{f_{AI}\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}}{f\left(\mathbb{E}_z(z^{\sigma-1})-\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}\right)}\right)^{\frac{1}{(\sigma-1)(1-\alpha)}}$, which is a constant, we have $s_a^{AI} = C \cdot s_p^{AI}$ from the condition of indifference about the expected profits between using AI and not using at the AI cutoff. Note that

$$h(s_p^{AI}) \equiv \int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f dF(s) + \int_{Cs_p^{AI}}^{+\infty} \left[\left(\frac{s}{Cs_p^{AI}}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f_{AI} dF(s)$$

is a continuous function of s_p^{AI} on the domain of $(0,+\infty)$ and is strictly decreasing in s_p^{AI}. What's more, $\lim_{s_p^{AI}\to 0} h(s_p^{AI}) = +\infty$ and $\lim_{s_p^{AI}\to\infty} h(s_p^{AI}) = 0$. We can then conclude that there exists a unique $s_p^{AI} > 0$ that satisfies the free entry condition $\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z(\pi^{AI}(s,z))dF(s) = f_e$. Then s_a^{AI} is also uniquely determined. The price index implies that zero expected profit condition at the cutoff satisfies. Finally, the number of entrants is given by the labor market clear condition

$$N_e^{AI} = \frac{1}{\sigma\left[f\int_{s_p^{AI}}^{+\infty}\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)} dF(s) + f_{AI}\int_{s_a^{AI}}^{+\infty}\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)} dF(s)\right]},$$

as we have solved for s_p^{AI} and s_a^{AI}.

To complete the proof for the existence of the equilibrium, we need to verify that the price index given by the zero expected profit condition is exactly what consumers face from

the demand side. To see this, we derive

$$
\begin{aligned}
P &= \left[\int_{\omega\in\Omega}\left(\frac{p(\omega)}{z(\omega)s(\omega)}\right)^{1-\sigma}d\omega\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e^{AI}\int_{s_p^{AI}}^{+\infty}\mathbb{E}_z\left(\frac{p(s,z)}{zs}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e^{AI}\int_{s_p^{AI}}^{s_a^{AI}}\mathbb{E}_z\left(\frac{\sigma}{\sigma-1}\frac{z^{-\frac{1}{\sigma}}}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}s^{\alpha-1}\right)^{1-\sigma}dF(s)+N_e^{AI}\int_{s_a^{AI}}^{+\infty}\mathbb{E}_z\left(\frac{\sigma}{\sigma-1}z^{-1}s^{\alpha-1}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= N_e^{AI\ \frac{1}{1-\sigma}}\frac{\sigma}{\sigma-1}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\int_{s_p^{AI}}^{s_a^{AI}}s^{(\sigma-1)(1-\alpha)}dF(s)+\mathbb{E}_z z^{\sigma-1}\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= N_e^{AI\ \frac{1}{1-\sigma}\frac{\sigma}{\sigma-1}}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\int_{s_p^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+\left[\mathbb{E}_z z^{\sigma-1}-\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\right]\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\left[N_e^{AI}\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\left(\int_{s_p^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+\frac{f_{AI}}{f}\left(\frac{s_p^{AI}}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)}\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)\right)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\frac{s_p^{AI\ (\sigma-1)(1-\alpha)}}{\sigma f}\right]^{\frac{1}{1-\sigma}} \\
&= \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{AI\ 1-\alpha}}\left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.
\end{aligned}
$$

This is exactly what we define as the price index. We complete the proof of the existence of the equilibrium. What's more, as the solutions for s_p^{AI} and s_a^{AI} are unique, the uniqueness of the equilibrium thus follows. Now we finish the proof for the existence and uniqueness of the equilibrium with AI.

B.2.5 Proof of Proposition 1

Note that s_p satisfies

$$\int_{s_p}^{+\infty}\left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)=f_e$$

and s_p^{AI} satisfies

$$\int_{s_p^{AI}}^{+\infty}\left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)+\int_{Cs_p^{AI}}^{+\infty}\left[\left(\frac{s}{Cs_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{AI}dF(s)=f_e.$$

As $\int_{Cs_p^{AI}}^{+\infty}\left[\left(\frac{s}{Cs_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{AI}dF(s)>0$, we have

$$\int_{s_p^{AI}}^{+\infty}\left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)<\int_{s_p}^{+\infty}\left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s).$$

Since $h(s_p)\equiv\int_{s_p}^{+\infty}\left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)$ strictly decreases in s_p, we have $s_p^{AI}>s_p$.

B.3 Open Economy Case with AI

B.3.1 Firms' Export Decision and the Sorting Pattern

For firms that do not adopt AI and serve only the domestic market, the expected profit is analogous to (2):

$$\mathbb{E}_z\Big(\pi^{AI}(s,z|n)\Big)=\left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{\sigma}\Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1}-f.$$

On the other hand, for firms that adopt AI while serving the domestic market only, the expected profit is given by:

$$\mathbb{E}_z\Big(\pi^{AI}(s,z|a)\Big)=\frac{\mathbb{E}_z\left(z^{\sigma-1}\right)}{\sigma^{\sigma}}\Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1}-f-f_{AI}.$$

For exporters employing AI

$$q_e^{AI}(s,z_d|e)=\left[\frac{\sigma-1}{\tau\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}(z_dP)^{\frac{\sigma-1}{\sigma}}\right]^{\sigma}$$

for each foreign market with consumer preference parameter z_d and the corresponding price is

$$p_e^{AI}(s,z_d|e)=\frac{\tau\sigma}{\sigma-1}s^{\alpha}.$$

The expected profit (excluding entry cost) is

$$\mathbb{E}_z(\pi^{AI}(s,z|e))=(1+n\tau^{1-\sigma})\frac{\mathbb{E}_z z^{\sigma-1}}{\sigma^{\sigma}}\Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1}-f-f_{AI}-nf_{exp}.$$

Firms will choose to export if and only if

$$\mathbb{E}_z\Big(\pi^{AI}(s,z|e)\Big)\geq\mathbb{E}_z\Big(\pi^{AI}(s,z|a)\Big),$$

which gives us the export cutoff s_e^{AI} as

$$s_e^{AI} = \left[\frac{\tau f_{exp}^{\frac{1}{\sigma-1}} \sigma^{\frac{\sigma}{\sigma-1}}}{(\sigma-1)P(\mathbb{E}_z z^{\sigma-1})^{\frac{1}{\sigma-1}}} \right]^{\frac{1}{1-\alpha}} = \tau^{\frac{1}{1-\alpha}} s_p^{AI} \left(\frac{f_{exp} \left[\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}} \right]^{\sigma}}{f \left[\mathbb{E}_z z^{\sigma-1} \right]} \right)^{\frac{1}{(\sigma-1)(1-\alpha)}}.$$

By Assumption 3, $s_e^{AI} > s_a^{AI}$. Note that if a firm chooses to export but to not use AI, the gain from exporting would be smaller. Thus, the cutoff for exporting would be even higher, which implies this type of firms won't exist in equilibrium. We then obtain the sorting pattern as $s_p^{AI} < s_a^{AI} < s_e^{AI}$.

B.3.2 Equilibrium

From the zero expected profit condition at the production cutoff and indifference conditions at the AI cutoff and export cutoff, the expected profit for firms with quality $s \in (s_p^{AI}, s_a^{AI})$ can be written as

$$\mathbb{E}_z(\pi^{AI}(s,z|n)) = \left[\left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f,$$

that for firms with quality $s \in [s_a^{AI}, s_e^{AI})$ can be written as

$$\mathbb{E}_z(\pi^{AI}(s,z|a)) = \left[\left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f + \left[\left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI},$$

and that for firms with quality $s \in [s_a^{AI}, s_e^{AI})$ can be written as

$$\mathbb{E}_z(\pi^{AI}(s,z|e)) = \left[\left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f + \left[\left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI} + n \left[\left(\frac{s}{s_e^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{exp}.$$

From free entry condition, we have

$$\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z \left(\pi^{AI}(s,z) \right) dF(s) = f_e,$$

i.e.,

$$\int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f dF(s) + \int_{s_a^{AI}}^{+\infty} \left[\left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI} dF(s) + n \int_{s_e^{AI}}^{+\infty} \left[\left(\frac{s}{s_e^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{exp} dF(s) = f_e.$$

With the relationship between s_p^{AI}, s_a^{AI} and s_e^{AI}, we can solve for s_p^{AI}, s_a^{AI} and s_e^{AI} from the equations. Note that s_a^{AI} and s_e^{AI} increase with s_p^{AI}.

We now define the equilibrium with access to AI technology in the open economy.

Definition 3 (Equilibrium with AI under Free Trade) *Given* $\{F(s), G(z), \sigma, f, f_x, \tau\}$, *an equilibrium consists of* $\{q^{AI}(s,z), p^{AI}(s,z), q_e^{AI}(s,z), p_e^{AI}(s,z), s_p^{AI}, s_a^{AI}, s_e^{AI}, N_e^{AI}, P^{AI}\}$ *such that*

(1) customers maximize their utility;

(2) firms maximize their (expected) profit, which implies $p_e^{AI}(s,z_d) = \frac{\tau\sigma}{\sigma-1}s^{\alpha}$ *and* $q_e^{AI}(s,z_d) = \left[\frac{\sigma-1}{\tau\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} z_d^{\sigma-1}$ *in the foreign market, and* $p^{AI}(s,z) = \frac{\sigma}{\sigma-1}s^{\alpha}$ *and* $q^{AI}(s,z) = \left[\frac{\sigma-1}{\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} z^{\sigma-1}$ *in the domestic market for firms with* $s \in [s_e^{AI}, +\infty)$, $p^{AI}(s,z) = \frac{\sigma}{\sigma-1}s^{\alpha}$ *and* $q^{AI}(s,z) = \left[\frac{\sigma-1}{\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\right]^{\sigma} z^{\sigma-1}$ *in the domestic market for firms with* $s \in [s_a^{AI}, s_e^{AI})$, *and* $p^{AI}(s,z) = \frac{\sigma}{\sigma-1}\frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}s^{\alpha}$ *and* $q^{AI}(s) = \left[\frac{\sigma-1}{\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}$ *for firms with* $s \in [s_p^{AI}, s_a^{AI})$;

(3) zero expected profit (excluding entry cost) at the production cutoff

$$\mathbb{E}_z(\pi^{AI}(s_p^{AI}, z|n)) = 0;$$

(4) indifference about the expected profit between using AI and not using at the AI cutoff

$$\mathbb{E}_z(\pi^{AI}(s_a^{AI}, z|n)) = \mathbb{E}_z(\pi^{AI}(s_a^{AI}, z|a));$$

(5) indifference about the expected profit between exporting and not exporting at the export cutoff

$$\mathbb{E}_z(\pi^{AI}(s_e^{AI}, z|a)) = \mathbb{E}_z(\pi^{AI}(s_e^{AI}, z|e));$$

(6) free entry condition is satisfied

$$\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z(\pi^{AI}(s,z))dF(s) = f_e;$$

(7) goods market clears;

(8) labor market clears

$$N_e^{AI}\left\{\int_{s_p^{AI}}^{+\infty}\left[(\sigma-1)\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}+1\right]fdF(s) + \int_{s_a^{AI}}^{+\infty}\left[(\sigma-1)\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)}+1\right]f_{AI}dF(s) + n\int_{s_e^{AI}}^{+\infty}\left[(\sigma-1)\left(\frac{s}{s_e^{AI}}\right)^{(\sigma-1)(1-\alpha)}+1\right]f_{exp}dF(s) + f_e\right\} = 1;$$

(9) price index is given by

$$P^{AI} = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{AI\ 1-\alpha}}\left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.$$

B.3.3 Proof of the Existence and Uniqueness of Equilibrium

The solution to consumers' utility maximization problem is given by the demand function. Firms' profit maximization problem is solved jointly by their pricing decision $p^{AI}(s,z)$, quantity decision $q^{AI}(s,z)$, and production decision, AI adoption decision, and export decision, i.e., the cutoff strategy s_p^{AI} with zero expected profit at the cutoff, s_a^{AI} with indifferent expected profit at the cutoff and s_e^{AI} with indifferent expected profit at the cutoff. Note that the price set by the firm exactly clears the goods market.

We then first show that there exists a unique set of s_p^{AI}, s_a^{AI} and s_e^{AI} that satisfies the free entry condition. By defining $C_1 \equiv \left(\frac{f_{AI}\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}}{f\left(\mathbb{E}_z(z^{\sigma-1})-\left[\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}\right)} \right)^{\frac{1}{(\sigma-1)(1-\alpha)}}$ and $C_2 \equiv \tau^{\frac{1}{1-\alpha}} \left(\frac{f_{exp}\left[\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}\right]^{\sigma}}{f\left[\mathbb{E}_z z^{\sigma-1}\right]} \right)^{\frac{1}{(\sigma-1)(1-\alpha)}}$, which are two constants, we have $s_a^{AI} = C_1 \cdot s_p^{AI}$ and $s_e^{AI} = C_2 \cdot s_p^{AI}$ from the conditions of indifference about the expected profits between using AI and not using at the AI cutoff, and between exporting and not exporting at the export cutoff. Note that

$$h(s_p^{AI}) \equiv \int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f dF(s) + \int_{C_1 s_p^{AI}}^{+\infty} \left[\left(\frac{s}{C_1 s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI} dF(s) + n \int_{C_2 s_p^{AI}}^{+\infty} \left[\left(\frac{s}{C_2 s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{exp} dF(s)$$

is a continuous function of s_p^{AI} on the domain of $(0, +\infty)$ and is strictly decreasing in s_p^{AI}. What's more, $\lim_{s_p^{AI} \to 0} h(s_p^{AI}) = +\infty$ and $\lim_{s_p^{AI} \to \infty} h(s_p^{AI}) = 0$. We can then conclude that there exists a unique $s_p^{AI} > 0$ that satisfies the free entry condition $\int_{s_p^{AI}}^{+\infty} \mathbb{E}_z(\pi^{AI}(s,z)) dF(s) = f_e$. Then s_a^{AI} and s_e^{AI} are also uniquely determined. The price index implies that zero expected profit condition at the cutoff satisfies. Finally, the number of entrants is given by the labor market clear condition

$$N_e^{AI} = \frac{1}{\left\{ \int_{s_p^{AI}}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f dF(s) + \int_{s_a^{AI}}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f_{AI} dF(s) + n \int_{s_e^{AI}}^{+\infty} \left[(\sigma-1) \left(\frac{s}{s_e^{AI}} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f_{exp} dF(s) + f_e \right\}}$$

as we have solved for s_p^{AI}, s_a^{AI}, and s_e^{AI}.

To complete the proof for the existence of the equilibrium, we need to verify that the price index given by the zero expected profit condition is exactly what consumers face from

the demand side. To see this, we derive

$$
\begin{aligned}
P &= \left[\int_{\omega\in\Omega}\left(\frac{p(\omega)}{z(\omega)s(\omega)}\right)^{1-\sigma}d\omega\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e^{AI}\int_{s_p^{AI}}^{+\infty}\mathbb{E}_z\left(\frac{p(s,z)}{zs}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \left[N_e^{AI}\int_{s_p^{AI}}^{s_a^{AI}}\mathbb{E}_z\left(\frac{\sigma}{\sigma-1}\frac{z^{-\frac{1}{\sigma}}}{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}s^{\alpha-1}\right)^{1-\sigma}dF(s)+N_e^{AI}\int_{s_a^{AI}}^{+\infty}\mathbb{E}_z\left(\frac{\sigma}{\sigma-1}z^{-1}s^{\alpha-1}\right)^{1-\sigma}dF(s)+nN_e^{AI}\int_{s_e^{AI}}^{+\infty}\mathbb{E}_z\left(\frac{\tau\sigma}{\sigma-1}z^{-1}s^{\alpha-1}\right)^{1-\sigma}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= N_e^{AI\ \frac{1}{1-\sigma}}\frac{\sigma}{\sigma-1}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\int_{s_p^{AI}}^{s_a^{AI}}s^{(\sigma-1)(1-\alpha)}dF(s)+\mathbb{E}_z z^{\sigma-1}\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+n\tau^{1-\sigma}\mathbb{E}_z z^{\sigma-1}\int_{s_e^{AI}}^{+\infty}s^{(\alpha-1)(1-\sigma)}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= N_e^{AI\ \frac{1}{1-\sigma}}\frac{\sigma}{\sigma-1}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\int_{s_p^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+\left(\mathbb{E}_z z^{\sigma-1}-\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\right)\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+n\tau^{1-\sigma}\mathbb{E}_z z^{\sigma-1}\int_{s_e^{AI}}^{+\infty}s^{(\alpha-1)(1-\sigma)}dF(s)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\left[N_e^{AI}\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\left(\int_{s_p^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+\frac{f_{AI}}{f}\left(\frac{s_p^{AI}}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)}\int_{s_a^{AI}}^{+\infty}s^{(\sigma-1)(1-\alpha)}dF(s)+\frac{f_{exp}}{f}\left(\frac{s_p^{AI}}{s_e^{AI}}\right)^{(\sigma-1)(1-\alpha)}\int_{s_e^{AI}}^{+\infty}s^{\sigma-1}(1-\alpha)dF(s)\right)\right]^{\frac{1}{1-\sigma}} \\
&= \frac{\sigma}{\sigma-1}\left[\left(\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right)^{\sigma}\frac{s_p^{AI\ (\sigma-1)(1-\alpha)}}{\sigma f}\right]^{\frac{1}{1-\sigma}} \\
&= \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{AI\ 1-\alpha}}\left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}.
\end{aligned}
$$

This is exactly what we define as the price index. We complete the proof of the existence of the equilibrium. What's more, as the solutions for s_p^{AI}, s_a^{AI} and s_e^{AI} are unique, the uniqueness of the equilibrium thus follows. Now we finish the proof for the existence and uniqueness of the equilibrium with AI in the open economy.

B.3.4 Benchmark: Open Economy Case without AI

In an open economy without AI, firms would choose between exporting and not exporting. If they do not export, the expected profit (excluding entry cost) for a firm with quality s would be

$$\mathbb{E}_z\Big(\pi(s,z|n)\Big)=\left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{\sigma}\Big[(\sigma-1)s^{1-\alpha}P\Big]^{\sigma-1}-f.$$

If they export, the quantity produced for each foreign market and the corresponding price would be

$$q_e(s,z_d|e)=\left[\frac{\sigma-1}{\tau\sigma}s^{\frac{\sigma-1}{\sigma}-\alpha}P^{\frac{\sigma-1}{\sigma}}\mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}$$

and

$$p_e(s, z_d|e) = \frac{\sigma}{\sigma - 1} \frac{z_d^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} \tau s^{\alpha}.$$

The expected profit (excluding entry cost) would be

$$\mathbb{E}_z\Big(\pi(s, z|e)\Big) = \Big(1 + n\tau^{1-\sigma}\Big) \left[\frac{\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)}{\sigma}\right]^{\sigma} \Big[(\sigma - 1)s^{1-\alpha}P\Big]^{\sigma-1} - f - nf_{exp}.$$

Firms will choose to export if and only if

$$\mathbb{E}_z\Big(\pi(s, z|e)\Big) \geq \mathbb{E}_z\Big(\pi(s, z|n)\Big),$$

which gives us the export cutoff s_e as

$$s_e = \left[\frac{\tau f_{exp}^{\frac{1}{\sigma-1}} \sigma^{\frac{\sigma}{\sigma-1}}}{(\sigma - 1)P\left[\mathbb{E}_z\left(z^{\frac{\sigma-1}{\sigma}}\right)\right]^{\frac{\sigma}{\sigma-1}}}\right]^{\frac{1}{1-\alpha}} = s_p \left[\frac{\tau^{\sigma-1} f_{exp}}{f}\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}.$$

By Assumption 3, $s_e > s_p$.

From free entry condition, we have

$$\int_{s_p}^{+\infty} \mathbb{E}_z\left(\pi(s, z)\right) dF(s) = f_e,$$

i.e.,

$$\int_{s_p}^{+\infty} \left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f dF(s) + n\int_{s_e}^{+\infty} \left[\left(\frac{s}{s_e}\right)^{(\sigma-1)(1-\alpha)} - 1\right] f_{exp} dF(s) = f_e.$$

Definition 4 (Equilibrium without AI under Free Trade) *Given* $\{F(s), G(z), \sigma, f, f_x, \tau\}$, *an equilibrium consists of* $\{q(s, z), p(s, z), q_e(s, z), p_e(s, z), s_p, s_e, N_e, P\}$ *such that*

(1) customers maximize their utility;

(2) firms maximize their (expected) profit, which implies $p_e(s, z_d) = \frac{\sigma}{\sigma-1} \frac{z_d^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} \tau s^{\alpha}$ *and* $q_e(s, z_d) = \left[\frac{\sigma-1}{\tau\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} P^{\frac{\sigma-1}{\sigma}} \mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}$ *for firms with* $s \in [s_e, +\infty)$, *in the foreign market, and* $p(s, z) = \frac{\sigma}{\sigma-1} \frac{z^{\frac{\sigma-1}{\sigma}}}{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}} s^{\alpha}$ *and* $q^{AI}(s, z) = \left[\frac{\sigma-1}{\sigma} s^{\frac{\sigma-1}{\sigma}-\alpha} P^{\frac{\sigma-1}{\sigma}} \mathbb{E}_z(z^{\frac{\sigma-1}{\sigma}})\right]^{\sigma}$ *for firms with* $s \in [s_p, +\infty)$ *in the domestic market;*

(3) zero expected profit (excluding entry cost) at the production cutoff

$$\mathbb{E}_z(\pi(s_p, z|n)) = 0;$$

(4) indifference about the expected profit between exporting and not exporting at the export cutoff

$$\mathbb{E}_z(\pi(s_e, z|n)) = \mathbb{E}_z(\pi(s_e, z|e));$$

(5) free entry condition is satisfied

$$\int_{s_p}^{+\infty} \mathbb{E}_z(\pi(s, z))dF(s) = f_e;$$

(6) goods market clears;

(7) labor market clears

$$N_e \left\{ \int_{s_p}^{+\infty} \left[(\sigma - 1) \left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f dF(s) + n \int_{s_e}^{+\infty} \left[(\sigma - 1) \left(\frac{s}{s_e} \right)^{(\sigma-1)(1-\alpha)} + 1 \right] f_{exp} dF(s) + f_e \right\} = 1;$$

(8) price index is given by

$$P = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma - 1)s_p^{1-\alpha}} \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma} \right]^{-\frac{\sigma}{\sigma-1}}.$$

The existence and uniqueness of the equilibrium can be proved *a la* the cases above.

B.3.5 Proof of Proposition 2

Note that s_p satisfies

$$\int_{s_p}^{+\infty} \left[\left(\frac{s}{s_p} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f dF(s) + n \int_{s_e}^{+\infty} \left[\left(\frac{s}{s_e} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{exp} dF(s) = f_e$$

and s_p^{AI} satisfies

$$\int_{s_p^{AI}}^{+\infty} \left[\left(\frac{s}{s_p^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f dF(s) + \int_{s_a^{AI}}^{+\infty} \left[\left(\frac{s}{s_a^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{AI} dF(s) + n \int_{s_e^{AI}}^{+\infty} \left[\left(\frac{s}{s_e^{AI}} \right)^{(\sigma-1)(1-\alpha)} - 1 \right] f_{exp} dF(s) = f_e.$$

With the definition of $C_2 \equiv \tau^{\frac{1}{1-\alpha}} \left(\frac{f_{exp} \left[\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}} \right]^{\sigma}}{f \left[\mathbb{E}_z z^{\sigma-1} \right]} \right)^{\frac{1}{(\sigma-1)(1-\alpha)}}$ and $C_3 \equiv \left[\frac{\tau^{\sigma-1} f_{exp}}{f} \right]^{\frac{1}{(\sigma-1)(1-\alpha)}}$, we have $s_e^{AI} = C_2 s_p^{AI}$ and $s_e = C_3 s_p$. Since $C_2 \leq C_3$, if $s_p \geq s_p^{AI}$, we would have $s_e \geq s_e^{AI}$.

This further implies

$$
\begin{aligned}
f_e = & \int_{s_p^{AI}}^{+\infty}\left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)+\int_{s_a^{AI}}^{+\infty}\left[\left(\frac{s}{s_a^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{AI}dF(s)+n\int_{s_e^{AI}}^{+\infty}\left[\left(\frac{s}{s_e^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{exp}dF(s) \\
> & \int_{s_p^{AI}}^{+\infty}\left[\left(\frac{s}{s_p^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)+n\int_{s_e^{AI}}^{+\infty}\left[\left(\frac{s}{s_e^{AI}}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{exp}dF(s) \\
\geq & \int_{s_p}^{+\infty}\left[\left(\frac{s}{s_p}\right)^{(\sigma-1)(1-\alpha)}-1\right]fdF(s)+n\int_{s_e}^{+\infty}\left[\left(\frac{s}{s_e}\right)^{(\sigma-1)(1-\alpha)}-1\right]f_{exp}dF(s) \\
= & f_e,
\end{aligned}
$$

which is impossible. Hence, we have $s_p < s_p^{AI}$, which is the first part of Proposition 2.

For exporting firms, the expectations of quantity sample variance across foreign markets with and without AI are [40]

$$
\begin{aligned}
\mathbb{E}(s^2(q)) &= \mathbb{E}\left(\frac{1}{n}\sum_{d=1}^{n}\left(q_d-\frac{\sum_{i=1}^{n}q_i}{n}\right)^2\right) \\
&= \frac{n-1}{n}Var(q)
\end{aligned}
$$

as random variables q_d are i.i.d in both cases. Without AI, since $q_d = q,\ \forall d$, this equation equals 0, and with AI, this equation is larger than zero. Hence, the quantity sample variance across foreign markets increases in expectation.

Similarly, the expectations of price sample variance across foreign markets with and without AI are

$$
\begin{aligned}
\mathbb{E}(s^2(p)) &= \mathbb{E}\left(\frac{1}{n}\sum_{d=1}^{n}\left(p_d-\frac{\sum_{i=1}^{n}p_i}{n}\right)^2\right) \\
&= \frac{n-1}{n}Var(p)
\end{aligned}
$$

as random variables p_d are i.i.d in both cases. Without AI, this equation is larger than zero, and with AI, since $p_d = p,\ \forall d$, this equation equals 0. Hence, the price sample variance across foreign markets decreases in expectation. We thus arrive at the second part of Proposition 2.

The price index, given by

$$
P = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{1-\alpha}}\left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}
$$

[40] Here we define s^2 with the denominator of n rather $n-1$ which is more often used. This constant multiplier would not impact our results at all as long as we use the same definition in both cases.

and

$$P^{AI} = \frac{f^{\frac{1}{\sigma-1}}}{(\sigma-1)s_p^{AI\ 1-\alpha}} \left[\frac{\mathbb{E}_z z^{\frac{\sigma-1}{\sigma}}}{\sigma}\right]^{-\frac{\sigma}{\sigma-1}}$$

respectively, becomes smaller with the availability of AI, since $s_p < s_p^{AI}$ as shown above. Thus, we complete the proof of Proposition 2.

C Anecdotal Evidence: Examples of AI adoptions

To give a concrete sense of how firms use AI to mitigate demand-side information frictions in export markets, we present illustrative job postings from our vacancy data and complementary examples from industry practice.

Job-level evidence from vacancies. Figure A.6 shows representative postings in which employers explicitly describe AI use in market-facing functions. Market intelligence and business analytics roles emphasize AI-enabled demand forecasting, customer segmentation, and competitive intelligence, often involving the construction of data pipelines and dashboards that translate customer signals into actionable market insights. Commercial strategy and sales-operations positions highlight predictive modeling to prioritize leads, target markets, and channels, and to support go-to-market decisions. Pricing and revenue-management roles stress algorithmic tools for price setting, performance monitoring, and rapid updating of pricing recommendations as market conditions evolve. These postings suggest a common theme: firms use AI in sales and marketing to turn past transactions and ongoing market feedback into sharper signals about demand. In our context, this is exactly the margin that matters—better demand prediction that guides pricing, destination entry and exit, and how shipments are allocated across export markets.

Business AI applications. Beyond our dataset, we collect actual business cases that illustrate how differentiated AI functions help exporters and multinationals understand markets.[41] Figure A.1 presents the news reported by *China Council For the Promotion of International Trade* on AI-driven marketing and intelligent exporting, highlighting the rapid

[41] AI has rapidly transformed sales and marketing practices at prominent companies worldwide. Leading technology firms like Amazon benefit substantially from deploying AI in recommendation engines, supply chain optimization, and sophisticated market analytics (Ruby Engine, 2024; AWS, 2025). Traditional retail and sales firms, such as Zara and Walmart, leverage AI to accurately forecast consumer demand trends, optimize inventory management, and deliver personalized customer experiences (Abdullah,2025; Walmart Global Tech, 2024). Specialized technology providers offer AI-driven digital transformation solutions, helping other companies integrate AI capabilities such as customer relationship management (CRM), predictive analytics, and automated business processes into their existing operations (See Salesforce for examples).

growth of AI-powered business analytics and marketing strategies in China's foreign trade sector. Compared to its well-documented production-side effects such as automation and cost reduction, AI's impact on the demand side remains an underexplored dimension. Empirical evidence from industry reports and business case studies reveals that firms integrating AI into market forecasting significantly outperform traditional forecasting methods, reducing prediction errors by approximately 20% to 50%, and consequently experiencing 10-15% gains in sales growth, 20-30% reduce in inventory costs, and improved market alignment.[42] As illustrated in Figure A.2, business-related adoption of AI can deliver substantial value, particularly in service operations, product development, and sales & marketing, underscoring AI's demand-side performances. We also showcase various examples of Chinese firms utilizing AI, including software as a service (SaaS), advanced planning and scheduling, and production chain optimizations.

E.g.1 "Currently, Bailian Intelligence has established a comprehensive and mature enterprise database along with an industry knowledge base matrix, demonstrating industry-leading AI application capabilities. Its product portfolio not only includes intelligent customer acquisition SaaS products but also features independently developed AI-driven applications for data security classification and grading, policy interpretation, and official document writing. By advancing in both customer acquisition scenarios and underlying AI technology applications, the company is pursuing a dual-track approach to innovation."

E.g.2 "ZhongAn Technology leverages artificial intelligence, big data, and cloud computing to drive innovation while integrating the advantages of the ZhongAn ecosystem. Validated through real-world business applications, the company has developed a value delivery system combining technology and services, centered around three standardized technology product series: business growth, business operations, and business infrastructure."

E.g.3 "Midea's digital supply chain platform enables real-time production planning for over 6,000 core upstream suppliers. Market order data from various sales channels converge onto Midea's digital platform, where upstream and downstream enterprises can share end-market demand information in real-time via the cloud. Leveraging the platform's intelligent scheduling algorithms, suppliers can develop master production plans

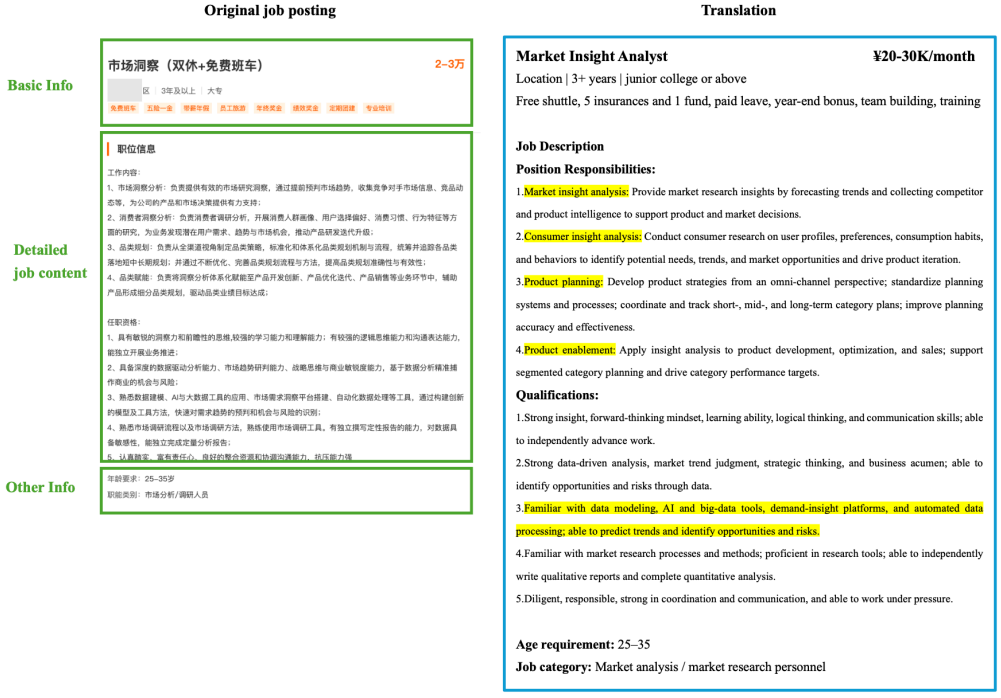
[42] McKinsey Report (2025) and Deloitte Insights (2025) provide comprehensive cross-nation analysis on how AI is reshaping business strategies and value creation. Ozturk (2024) and the article in Global Trade (2024) compiles detailed evidence of AI-powered forecasting and trade performance across regions and sectors, including retail supply chains, food & beverage, fast fashion, and global logistics.

and refine them into detailed sub-production plans down to individual workstations and equipment. This flexible production approach enhances the responsiveness and coordination across the entire industrial chain, improving order fulfillment rates and better adapting to the increasingly fragmented and dynamic demands of downstream markets."

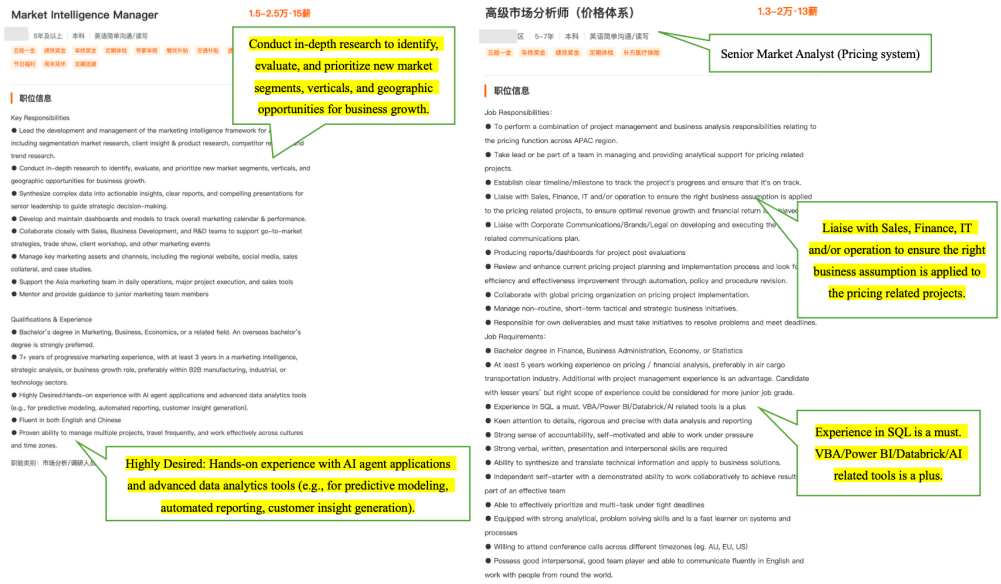
In summary, the real-world cases – from AI-powered personalized recommendations to demand forecasting and marketing – demonstrate how businesses are deploying AI to sharpen their knowledge of consumer preferences. This market-learning function of AI distinguishes it from traditional automation technologies. While production-oriented AI might simply reduce costs uniformly across markets, sales-oriented AI allows firms to tailor their strategies to specific market segments by lowering information frictions. Our theoretical model formalizes how this capability reshapes export behavior and market outcomes.

Figure A.6: Illustrative Examples of AI-Sales related Job Postings

(a) A fully translated and structured job posting

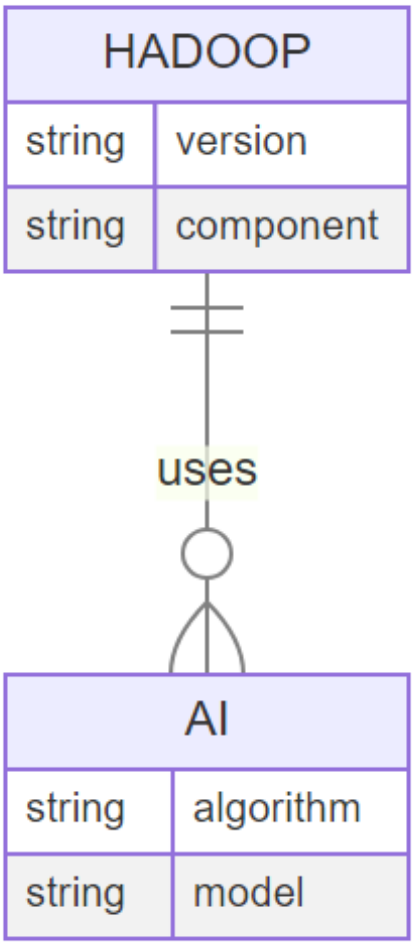


(b) Examples with highlights on AI-sales related content



Note: These examples are actual job postings we draw from our data to illustrate the standardized fields, textual structure, and AI-specific content. The location information is masked on purpose.

Figure A.7: Hadoop as Data Infrastructure for AI Workflows



Notes: The figure illustrates why some keywords in our AI list refer to the AI *stack* rather than the term "AI" itself. Hadoop provides distributed storage and data-processing infrastructure that supports data preparation and large-scale computation, which commonly underpin machine-learning model training and deployment. Source: https://blog.51cto.com/u_16175440/12826253.

D Robustness Checks

Table A.4: Robustness Checks: Impact of AI Adoption on Price and Quantity Dispersion

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(q)_{ift}$	(3) $\ln Var(p)_{ift}$	(4) $\ln Var(q)_{ift}$	(5) $\ln Var(p)_{ift}$	(6) $\ln Var(q)_{ift}$	(7) $\ln Var(p)_{ift}$	(8) $\ln Var(q)_{ift}$
	Panel A: Excluding National Municipalities				*Panel B: Winsorizing Extreme Values*			
AI^S_{ift}	-2.098***	0.282*	-1.726***	0.259*	-2.447***	0.252*	-2.135***	0.294**
	(0.419)	(0.146)	(0.414)	(0.126)	(0.500)	(0.124)	(0.479)	(0.137)
Observations	77,756	76,950	73,520	72,798	113,250	111,820	106,510	105,210
Adjusted R^2	0.767	0.899	0.776	0.900	0.781	0.898	0.790	0.900
	Panel C: Assigning Zero to Missing Vacancy Values				*Panel D: Rescaling Product Shares Relative to Core Product*			
AI^S_{ift}	-2.637***	0.258**	-2.346***	0.285**	-1.016***	0.227***	-0.823***	0.190***
	(0.508)	(0.111)	(0.476)	(0.112)	(0.213)	(0.033)	(0.157)	(0.043)
Observations	182,224	179,944	169,948	167,940	114,466	113,018	107,688	106,378
Adjusted R^2	0.780	0.899	0.790	0.902	0.778	0.900	0.788	0.902
	Panel E: Using Core Product Dummy				*Panel F: Using AI-sales Hiring Intensity*			
AI^S_{ift}	-1.379***	0.191**	-1.157***	0.180**	-3.795***	3.742**	-4.810***	3.291**
	(0.390)	(0.067)	(0.308)	(0.074)	(1.215)	(1.384)	(1.252)	(1.497)
Observations	113,250	111,820	106,510	105,210	51,980	51,300	55,294	54,560
Adjusted R^2	0.778	0.900	0.788	0.902	0.772	0.900	0.762	0.898
Fixed Effects:								
Firm-HS 8-digit	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y	Y	Y	Y	Y

Note: Panel A excludes firms located in the four municipalities directly administered by the central government of China (Beijing, Tianjin, Shanghai, and Chongqing). Panel B winsorizes extreme values of the AI variable at the top and bottom 1%. Panel C treats missing vacancy postings as zero, assuming that firms did not recruit sales-AI workers in years when they do not appear in the vacancy dataset. Panel D replaces the continuous product share measure with a binary indicator for the core product —defined as the product with the highest export share within each firm-year. Panel E interacts AI adoption with the core product dummy, rather than the continuous share variable. Panel F interacts share with firm-level AI-sales hiring intensity, rather the dummy variable. Columns (1)(2)(5)(6) report results on dispersion across transactions and other columns report results on dispersion across markets. Robust standard errors, clustered at the industry level, are reported in parentheses. *** p <0.01, ** p <0.05, * p <0.1.

Table A.5: Robustness Checks: Comparing AI-sales Hiring with Traditional Sales Hiring

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(q)_{ift}$	(4) $\ln Var(q)_{ift}$
Panel A: Price and Quantity Dispersion Across Destinations				
AI^S_{ift}	-4.079***	-4.043***	0.333**	0.355**
	(0.429)	(0.442)	(0.129)	(0.128)
$Sales^{Non-AI}_{ift}$	-2.182***	-2.131***	0.0764	0.103**
	(0.215)	(0.213)	(0.0465)	(0.0469)
$\ln(\#\text{ of Destinations})_{ift}$		0.758***		0.413***
		(0.015)		(0.019)
Observations	106,510	106,510	105,210	105,210
Adjusted R^2	0.789	0.793	0.902	0.903
Panel B: Price and Quantity Dispersion Across All Transactions				
AI^S_{ift}	-4.583***	-4.521***	0.323**	0.356**
	(0.425)	(0.438)	(0.124)	(0.123)
$Sales^{Non-AI}_{ift}$	-2.465***	-2.412***	0.126***	0.152***
	(0.224)	(0.219)	(0.0426)	(0.0442)
$\ln(\#\text{ of Transactions})_{ift}$		0.792***		0.393***
		(0.011)		(0.010)
Observations	113,250	113,250	111,820	111,820
Adjusted R^2	0.779	0.784	0.900	0.901
Fixed Effects:				
Firm-HS 8-digit	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y

Note: We compare the effects of AI-sales hiring with traditional sales hiring. AI^S_{ift} is as defined in equation (8), and $Sales^{Non-AI}_{ift}$ is similarly defined, except that AI^S_{ft} in the equation is replaced with a dummy indicating firm-level traditional sales hiring up to year t. Panel A reports dispersion in prices and quantities across destinations within a firm–product–year, calculated as the variance in average unit prices and quantities sold across all destination markets that the firm serves with that product in a given year, as defined in equation (7). Panel B measures dispersion across individual transactions within each firm–product–year, using the full set of shipment-level trade flows, including repeated transactions to the same destination. $\ln(\#\text{ of Destinations})_{ift}$ and $\ln(\#\text{ of Transactions})_{ift}$ denote the log number of destinations and total transactions, respectively, for firm f selling product i in year t. Robust standard errors are clustered at the industry level and reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.6: Robustness Checks: Excluding Chinese Translation of Two AI Keywords

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(q)_{ift}$	(3) $\ln Var(p)_{ift}$	(4) $\ln Var(q)_{ift}$	(5) N_{ift}^{entry}	(6) N_{ift}^{exit}
AI_{ift}^{S}	-2.187***	0.309**	-2.508***	0.266**	0.971**	0.301**
	(0.497)	(0.137)	(0.514)	(0.122)	(0.358)	(0.116)
Observations	106,510	105,210	113,250	111,820	214,742	214,742
Adjusted R^2	0.788	0.902	0.778	0.900	0.815	0.896
Fixed Effects:						
Firm-HS 8-digit	Y	Y	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y	Y	Y

Note: Column (1) and (2) report dispersion in prices and quantities across destinations within a firm–product–year, calculated as the variance in average unit prices and quantities sold across all destination markets that the firm serves with that product in a given year, as defined in equation (7). Column (3) and (4) measure dispersion across individual transactions within each firm–product–year, using the full set of shipment-level trade flows, including repeated transactions to the same destination. In column (5) and (6), variables N_{ift}^{entry} and N_{ift}^{exit} denote the number of destination markets newly entered and exited, respectively, for each firm-product-year observation. Robust standard errors, clustered at the industry level, are reported in parentheses. *** p <0.01, ** p <0.05, * p <0.1.

Table A.7: Robustness Checks: the Distinct Role of Sales AI Versus Other Technologies

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(q)_{ift}$	(4) $\ln Var(q)_{ift}$	(5) $\ln Var(p)_{ift}$	(6) $\ln Var(p)_{ift}$	(7) $\ln Var(q)_{ift}$	(8) $\ln Var(q)_{ift}$
	Panel A: Non-Sales AI Effects on Market Dispersion				*Panel B: Non-Sales AI on Transaction Dispersion*			
AI^{NS}_{ift}	-0.993***	-0.953***	0.104	0.122	-1.033***	-0.997***	0.149	0.163
	(0.0577)	(0.056)	(0.107)	(0.108)	(0.141)	(0.132)	(0.119)	(0.118)
$\ln(\#\text{ of Destinations})_{ift}$		0.759***		0.413***				
		(0.014)		(0.019)				
$\ln(\#\text{ of Transactions})_{ift}$						0.794***		0.393***
						(0.011)		(0.010)
Observations	106,510	106,510	105,210	105,210	113,250	113,250	111,820	111,820
Adjusted R^2	0.788	0.793	0.902	0.903	0.778	0.783	0.900	0.901
	Panel C: Combining Sales and Non-Sales AI				*Panel D: AI Compared to Other Data Technologies*			
AI^{Pool}_{ift}	-2.044***	-2.008***	0.249**	0.266**	-1.613***	-1.572***	0.234*	0.253**
	(0.166)	(0.180)	(0.112)	(0.108)	(0.161)	(0.179)	(0.121)	(0.117)
$Data_{ift}$					-1.596***	-1.613***	0.0565	0.0496
					(0.122)	(0.132)	(0.057)	(0.054)
$\ln(\#\text{ of Destinations})_{ift}$		0.758***		0.413***		0.759***		0.413***
		(0.014)		(0.019)		(0.014)		(0.019)
Observations	106,510	106,510	105,210	105,210	106,510	106,510	105,210	105,210
Adjusted R^2	0.788	0.793	0.902	0.903	0.789	0.793	0.902	0.903
	Panel E: Sales AI with Controls for Other Technologies							
AI^{S}_{ift}	-2.588***	-2.572***	0.412**	0.420**				
	(0.584)	(0.627)	(0.172)	(0.157)				
AI^{NS}_{ift}	-1.414***	-1.368***	0.198	0.219*				
	(0.106)	(0.121)	(0.125)	(0.122)				
$Data_{ift}$	-1.510***	-1.525***	0.041	0.035				
	(0.122)	(0.132)	(0.053)	(0.051)				
$\ln(\#\text{ of Destinations})_{ift}$		0.759***		0.413***				
		(0.014)		(0.019)				
Observations	106,510	106,510	105,210	105,210				
Adjusted R^2	0.789	0.793	0.902	0.903				
Fixed Effects:								
Firm-HS 8-digit	Y	Y	Y	Y	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y	Y	Y	Y	Y

Note: Panel A reports results on the dispersion of prices and quantities across destinations within a firm–product–year, calculated as the variance in average unit prices and quantities across all destination markets served, as defined in equation (7), using non-sales AI as the independent variable. Panel B measures dispersion across individual transactions within each firm–product–year, based on shipment-level trade flows, including repeated transactions to the same destination, again using non-sales AI as the independent variable. Panel C instead uses overall AI adoption—regardless of application—as the independent variable. Panel D adds general data technology adoption as an additional independent variable. Panel E includes all three measures—sales AI, non-sales AI, and general data technology—simultaneously in the regression. $\ln(\#\text{ of Destinations})_{ift}$ and $\ln(\#\text{ of Transactions})_{ift}$ denote the log number of destinations and total transactions, respectively, for firm f selling product i in year t. Robust standard errors, clustered at the industry level, are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.8: Robustness Checks: Differencing Out Fixed Effects when Constructing Dispersion Variables

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(q)_{ift}$	(4) $\ln Var(q)_{ift}$
Panel A: Differencing Out HS-6-digit-Market and HS-6-digit-Year Fixed Effects				
AI^S_{ift}	-0.672***	-0.683***	0.207**	0.199**
	(0.166)	(0.182)	(0.0919)	(0.0733)
$\ln(\#\text{ of Destinations})_{ift}$		0.619***		0.460***
		(0.013)		(0.021)
Observations	105,784	105,784	105,782	105,782
Adjusted R^2	0.870	0.875	0.908	0.911
Panel B: Differencing Out HS-6-digit-Market and HS-8-digit-Year Fixed Effects				
AI^S_{ift}	-0.680***	-0.690***	0.215***	0.207***
	(0.172)	(0.188)	(0.0709)	(0.0533)
$\ln(\#\text{ of Destinations})_{ift}$		0.609***		0.452***
		(0.017)		(0.022)
Observations	105,784	105,784	105,782	105,782
Adjusted R^2	0.871	0.876	0.911	0.914
Fixed Effects:				
Firm-HS 8-digit	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y

Note: We difference out fixed effects when constructing the dependent variables, i.e., price disperion and quantity dispersion. Panel A differences out HS-6-digit-Market fixed effects and HS-6-digit-Year fixed effects, while Panel B differences out HS-6-digit-Market fixed effects and HS-8-digit-Year fixed effects. $\ln(\#\text{ of Destinations})_{ift}$ denotes the log number of destinations for firm f selling product i in year t. Robust standard errors are clustered at the industry level and reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A.9: Robustness Checks: Identifying Sales AI Adoption via Occupation-Based Vacancies

Dep var.	(1) $\ln Var(p)_{ift}$	(2) $\ln Var(p)_{ift}$	(3) $\ln Var(p)_{ift}$	(4) $\ln Var(p)_{ift}$
Panel A: Price Dispersion Across Destinations				
AI^{S}_{ift}	-2.018***	-2.007***		
	(0.377)	(0.396)		
AI^{NS}_{ift}			-1.819***	-1.831***
			(0.186)	(0.195)
$\ln(\#\text{ of Destinations})_{ift}$		0.765***		0.765***
		(0.015)		(0.015)
Observations	105,602	105,602	105,602	105,602
Adjusted R^2	0.789	0.793	0.789	0.793
Panel B: Quantity Dispersion Across Destinations				
AI^{S}_{ift}	0.133**	0.136***		
	(0.046)	(0.044)		
AI^{NS}_{ift}			0.0268	0.0202
			(0.118)	(0.117)
$\ln(\#\text{ of Destinations})_{ift}$		0.412***		0.412***
		(0.020)		(0.020)
Observations	104,302	104,302	104,302	104,302
Adjusted R^2	0.902	0.903	0.902	0.903
Fixed Effects:				
Firm-HS 8-digit	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y

Note: Panel A reports results for price dispersion, calculated as the variance in average unit prices across all destination markets served by a firm-product-year, following equation (7). Panel B reports results for quantity dispersion, calculated analogously as the variance in quantities sold across destinations. In both panels, AI adoption is classified based on SOC occupational codes. $\ln(\#\text{ of Destinations})_{ift}$ denotes the logarithm of the number of destinations for firm f selling product i in year t. Robust standard errors, clustered at the industry level, are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

D.1 Robustness Checks: AI Adoption and Firm Productivity

Estimating Firm Productivity. We estimate productivity by specifying the following production function. Let y_{it} denote the logarithm of output for firm i in year t, and assume the production function takes the following form:

$$\begin{aligned} y_{it} &= \beta_0 + \beta_l l_{it} + \beta_k k_{it} + \omega_{it} + \varepsilon_{it} \\ &= \beta_l l_{it} + \phi(k_{it}) + \varepsilon_{it} \end{aligned}$$

where k_{it} and l_{it} denote the logarithms of capital and labor inputs, respectively, and $\phi(k_{it}) = \beta_0 + \beta_k k_{it} + \omega(k_{it})$ captures the flexible component of the production function. The term ω_{it} represents firm-specific productivity, which is the object of interest. To estimate it, we approximate $\phi(k_{it})$ using a third-order polynomial in k_{it} and obtain a consistent estimate of the output elasticity of labor, $\hat{\beta}_l$, by estimating the following OLS regression:

$$y_{it} = \delta_0 + \beta_l l_{it} + \sum_{j=1}^{3} \delta_j k_{it}^j + \varepsilon_{it}.$$

We can thus obtain the consistent estimate of ϕ_{it} and $\hat{\varepsilon_{it}}$, which is given by

$$\hat{\phi}_{it} = \hat{y}_{it} - \hat{\beta}_l l_{it} = \hat{\delta}_0 + \sum_{j=1}^{3} \hat{\delta}_j k_{it}^j.$$

Then we estimate β_0 and β_k using the GMM method. Specifically, for each set of candidate values of β_0^* and β_k^*, we compute the prediction for ω_{it} as $\hat{\omega}_{it} = \hat{\phi}_{it} - \beta_0^* - \beta_k^* k_{it}$. Assume the motion of ω satisfies $\omega_{it} = \gamma_1 \omega_{it-1} + e_{it}$. Then by regressing ω_{it} on ω_{it-1}, we could obtain the estimation for $\mathbb{E}(\omega_{it}|\omega_{it-1}) = \hat{\gamma}_1 \omega_{it-1}$ and also the estimation for e_{it}. We then make use of the moment condition $\mathbb{E}((\varepsilon_{it} + e_{it})|i_{it}) = 0$, where i_{it} is the investment of firm i in year t, and search for β_0 and β_k to minimize $(\hat{\varepsilon}_{it} + \hat{e}_{it})i_{it}$.

As a robustness check, we also construct the moment condition using information on intermediate inputs, imposing $\mathbb{E}[(\varepsilon_{it} + e_{it}) \mid m_{it}] = 0$, where m_{it} denotes the logarithm of intermediate input usage. We then estimate the parameters β_0 and β_k accordingly. Based on these estimates, firm-level productivity is inferred as the residual. To mitigate the influence of outliers, we trim the estimated productivity distribution by 1% at both tails. Columns (1)–(3) of Table A.10 report results using the moment condition based on intermediate inputs, while Columns (4)–(6) present results using the investment-based moment condition.

Impact of AI Use On Firm Productivity. Column (1) uses overall AI adoption as the explanatory variable. While the coefficient is positive, it is not statistically significant,

Table A.10: Effects of AI Use on Firm Productivity

Dep var.	(1) $\ln TFP_{ft}$	(2) $\ln TFP_{ft}$	(3) $\ln TFP_{ft}$	(4) $\ln TFP_{ft}$	(5) $\ln TFP_{ft}$	(6) $\ln TFP_{ft}$
AI_{ft}^{Pool}	0.0376			0.0353		
	(0.0262)			(0.0259)		
AI_{ft}^{S}		-0.0177			-0.0295	
		(0.0415)			(0.0426)	
AI_{ft}^{NS}			0.0409**			0.0439**
			(0.0188)			(0.0185)
Observations	36,900	36,900	36,900	36,899	36,899	36,899
Adjusted R^2	0.664	0.664	0.664	0.647	0.647	0.647
Fixed Effects:						
Firm FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y

Note: The dependent variable is the total factor productivity of firm f in year t. Column (1) - (3) use productivity estimated from the moment condition based on intermediate inputs. Column (4) - (6) use productivity estimated form the moment condition based on investment. AI_{ft}^{Pool}, AI_{ft}^{S} and AI_{ft}^{NS} represent overall AI use, sales-related AI use, and non-sales AI use, respectively. Robust standard errors, clustered at the industry level, are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

suggesting no systematic link with productivity. Column (2) focuses on sales-related AI, and similarly finds an insignificant, slightly negative effect—consistent with its role in improving demand prediction and reducing market frictions rather than enhancing production efficiency. In contrast, Column (3) shows that non-sales AI—covering production, logistics, and management—is significantly associated with higher productivity. Columns (4)–(6) confirm these findings using an alternative estimation approach. Overall, the results underscore that sales-related AI improves market allocation rather than total factor productivity, and that distinguishing between functional types of AI adoption is essential for understanding firm-level impacts—an aspect often overlooked in existing research.

D.2 Effects on Average Export Prices and Quantities

The main text uses dispersion in prices and quantities across destinations to proxy for demand-side information frictions. Here we examine a complementary set of outcomes: average export prices and quantities at the firm–product level. Unlike dispersion, the model does not yield a sharp sign prediction for these averages. Better demand information induces firms to reallocate shipments toward destinations with stronger demand and away from weaker ones, and it also triggers market turnover through entry into attractive destinations and exit from unattractive ones (Table 5). Both channels can move averages, but not necessarily in the same direction.

To gauge the net effect in the data, we estimate regressions analogous to the baseline specification, replacing the dependent variable with the log average unit price and the log

total export quantity for each firm–product–year. Table A.11 reports the results. Greater exposure to sales AI is associated with lower average prices and a modest increase in exported quantities, with the quantity effect estimated less precisely.

Table A.11: Average Export Prices and Quantities

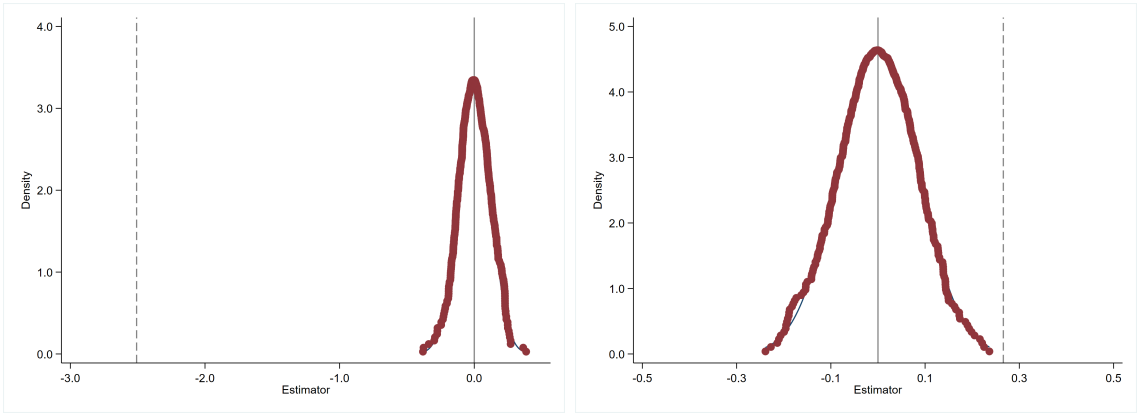
Dep var.	(1) $\ln(p)_{ift}$	(2) $\ln(q)_{ift}$	(3) $\ln(p)_{ift}$	(4) $\ln(q)_{ift}$
AI^S_{ift}	-0.541**	0.298*	-0.539**	0.249
	(0.192)	(0.170)	(0.192)	(0.167)
$\ln(\#\text{ of Destinations})_{ift}$			-0.0412***	1.095***
			(0.0102)	(0.0254)
Observations	214,742	214,742	214,742	214,742
Adjusted R^2	0.915	0.906	0.915	0.916
Fixed Effects:				
Firm-HS 8-digit	Y	Y	Y	Y
Firm-Year	Y	Y	Y	Y
HS 8-digit-Year	Y	Y	Y	Y

Note: The dependent variables are the log average unit price and log total export quantity at the firm–product–year level. Columns (3) and (4) control for the log number of destinations served by firm f for product i in year t. Robust standard errors clustered at the industry level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

These average effects also help discipline a pure "demand expansion" interpretation. If sales AI simply raised demand uniformly for the products and firms that adopt it, one would expect both prices and quantities to rise on average. Instead, we find lower average prices alongside (at most) a modest increase in quantities, consistent with improved targeting and market selection rather than a broad outward shift in demand. This interpretation is also in line with Table 5, where AI is associated with both more market entry and more market exit—a pattern that is difficult to reconcile with uniform demand growth, which would predict entry without a corresponding increase in exits.

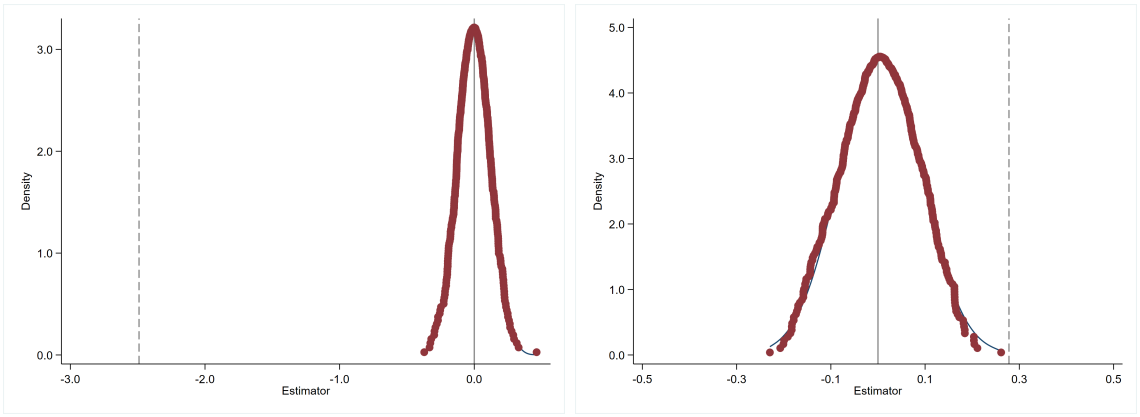
D.3 Robustness Checks: Figure of Placebo Tests

Figure A.8: Placebo Test for Estimated Effects of AI Adoption



Price Dispersion Quantity Dispersion

Panel A: Transaction-Level Dispersion



Price Dispersion Quantity Dispersion

Panel B: Destination-Level Dispersion

Note: In each iteration, we randomly reassign the AI adoption intensity across firms and re-estimate the baseline regression using the permuted treatment variable. We repeat this procedure 500 times and record the estimated coefficients on the placebo variable. The histograms plot the distribution of these coefficients for price dispersion (left) and quantity dispersion (right), separately for transaction-level data (Panel A) and destination-level data (Panel B). The vertical dashed lines indicate the coefficients obtained using the actual AI adoption intensity, showing that the baseline estimates lie well outside the placebo distribution.