The Pre-ChatGPT AI Era and Firm Sales: Domestic versus International Margins*

Suleyman Faruk Gozen† Pinar Gunes‡ Mehmet Furkan Karaca§

February 2026

Abstract

We study how artificial intelligence (AI) affects firms' sales in the pre-ChatGPT era. Exploiting the 2015 release of Google TensorFlow as a plausibly exogenous shock to AI-worker effectiveness, we implement a difference-in-differences design that leverages firm-level variation in pre-shock AI exposure. We show that AI adoption generates large and persistent increases in total, domestic, and foreign sales, but only among large firms; smaller firms exhibit no responses. Mechanism evidence suggests that while AI exposure raises operating efficiency and productivity for large firms, firms with financial constraints and limited intangible capital are less able to translate AI into sales gains.

Keywords: Artificial Intelligence, Domestic and Foreign Sales, Firm Size Distribution

JEL Codes: D22, F14, F16, F23, L11, L25, O33

*The previously circulated title was *"Artificial Intelligence, Trade, and Firm Dynamics"*. We thank Arpad Abraham, Emrehan Aktug, Seyit Mumin Cilasun, Banu Demir, Hakan Erdol, Michael Gasiorek, Robin Kaiji Gong, Ruveyda Gozen, Giammario Impullitti, Ayse Imrohoroglu, Paul Ko, Leandro de Magalhaes, Sergey Nigai, Zahra Siddique, Ruiqi Sun, Sami Stouli, Christine Valente, Hakki Yazici, Kamil Yilmaz, Mehmet Yorukoglu, Maurizio Zanardi, Yanos Zylberberg, and Tom Zylkin as well as workshop/conference participants at Bristol Research Day, Centre for Inclusive Trade Policy (CITP) Academic Conference (University of Strathclyde), IP Rights, Technology, and International Trade Research Workshop (University of Nottingham), Spring 2025 Midwest International Trade & Theory Conference (Virginia Tech), Spring Meeting of Turkish Economic Association 2025 (Istanbul Bilgi University), UK Department for Science, Innovation and Technology (DSIT) Meeting, UK Trade Policy Observatory (UKTPO) 7^{th} Annual Conference (Chatham House), and XVII. Winter Workshop in Economics at Koç University for helpful comments and suggestions. Computer programs to replicate the empirical analyses are available from the authors. All errors are our own.

†School of Economics, University of Bristol. Email: suleyman.gozen@bristol.ac.uk.

‡Business School, University of Sussex. Email: p.gunes@sussex.ac.uk.

§Essex Business School, University of Essex. Email: m.f.karaca@essex.ac.uk.

1 Introduction

Artificial intelligence (AI) is increasingly shaping how firms organize production, process information, and compete in domestic and international markets. While a growing literature documents the effects of AI on innovation, productivity, and labor demand (e.g. Agrawal et al., 2019, Babina et al., 2024, Cockburn et al., 2019, Hirvonen et al., 2022, Rock, 2021), far less is known about how AI adoption translates into firms' commercial performance across borders (e.g. Brynjolfsson et al., 2019, Chander, 2021, Ferencz et al., 2022, Goldfarb and Trefler, 2018, Lo and Lee, 2024, Xu and Tian, 2023) except for a few number of studies which focus on how AI shapes international trade patterns (e.g. Jakubik et al., 2025, Mishra et al., 2023, World Trade Organization, 2024, 2025). In particular, whether AI enables firms to expand internationally, and through which margins such expansion occurs, remains an open question. This gap is notable given the central role of firm heterogeneity, scale, and intangible assets in modern theories of international trade.

Existing research on technology and trade has largely focused on earlier waves of digitalization, such as information and communication technologies or automation, and their implications for export participation and productivity. AI, however, differs fundamentally from these technologies in its reliance on data, intangible capital, and organizational complementarities. These features suggest that AI may disproportionately benefit firms with sufficient scale and intangible capacity, precisely the firms most likely to engage in international activity. Yet causal evidence on how AI adoption affects firms' domestic and foreign sales margins remains scarce.

This paper provides new causal evidence on how AI adoption affects firms' domestic and international sales in the pre-ChatGPT era. Studying this period allows us to isolate the effects of AI adoption before the diffusion of large language models and generative AI, which introduced broad and simultaneous changes in firms' production, organization, and decision-making processes. We exploit the 2015 release of Google TensorFlow, a widely adopted open-source platform that sharply reduced the cost and complexity of implementing machine-learning algorithms. TensorFlow disproportionately increased the effectiveness of pre-existing AI-skilled labor, enabling firms with prior AI capabili-

ties to scale AI usage rapidly while leaving less-prepared firms largely unaffected. Using cross-firm variation in pre-2015 AI exposure, we implement a difference-in-differences framework to identify the causal impact of AI on firms' sales outcomes.

Our central finding is that the commercial gains from AI adoption are highly uneven across firms and strongly shaped by firm size. Following the release of TensorFlow, large firms with high pre-treatment AI intensity experience sizable and persistent increases in both total and foreign sales. These effects emerge quickly after the shock and remain stable over time, suggesting that AI enables large firms to scale their operations and expand their commercial footprint abroad. In contrast, smaller firms show no statistically meaningful response along any sales margin, indicating that AI adoption alone is insufficient to overcome the constraints associated with limited scale and organizational capacity. Importantly, we find little evidence of systematic effects on export sales. This divergence between foreign sales and exports implies that AI-driven international expansion operates primarily through foreign market activity such as the expansion of foreign affiliates or the strengthening of multinational production networks rather than through increased cross-border exports from the home country. Taken together, these patterns suggest that AI reinforces existing firm heterogeneity in international markets by disproportionately benefiting firms that are already positioned to operate at scale across borders.

Mechanism evidence further shows that large firms experience clear and sustained improvements in both labor productivity and cost efficiency after 2015, whereas small firms exhibit no comparable changes. For large firms, the post-2015 increases in productivity reflect their ability to reorganize production processes, integrate AI tools into existing workflows, and extract greater value from their accumulated intangible assets such as software capital, proprietary data, and organizational know-how. The improvements in cost efficiency indicate that AI does not simply raise output, but also enables more efficient deployment of inputs by optimizing scheduling, reducing operational frictions, and enhancing decision-making quality. These patterns are absent for smaller firms, which suggests that AI adoption requires a threshold level of scale and complementary organizational resources. In effect, AI acts as a force multiplier for firms that already possess the structural capacity to absorb and implement it, while firms below this threshold remain

unable to convert AI exposure into efficiency gains.

Additional heterogeneity analysis reveals that intangible capital amplifies the post-2015 gains from AI, and this amplification is substantially more pronounced for large firms. Firms with high intangible intensity derive markedly stronger improvements in sales, suggesting that the complementarities between AI technologies and intangible organizational assets scale disproportionately with firm size. In contrast, financial constraints generally weaken the post-2015 gains from AI exposure; however, this negative effect is considerably mitigated among large firms. Their greater internal resources and borrowing capacity allow them to offset the limiting effects of financial frictions, enabling them to translate AI exposure into commercial gains even when financially constrained.

To assess the robustness of our findings, we conduct a wide range of complementary checks. First, we implement placebo tests and find no evidence of differential pre-trends. Second, we show that our results are robust to alternative measures of AI exposure, including a narrower definition of AI-related workers. Third, we re-estimate our difference-in-differences specifications using propensity-score matching to address potential selection into AI adoption and obtain quantitatively similar results. We further demonstrate that our findings are robust to the inclusion of state-level and year-level linear trends, as well as to controlling for other firm-level technologies such as Information Technology (IT) and data analytics. Finally, our results remain unchanged when excluding technology-intensive and healthcare industries. Taken together, these exercises confirm that the estimated effects reflect a causal impact of AI adoption following the release of TensorFlow.

Our paper contributes to the several strands of the existing literature. First, it contributes to the literature on new technologies' impact on sales components by providing causal evidence on the impact of AI on firm-level domestic and foreign sales. To the best of our knowledge, this is the first paper that empirically investigates the impact of AI adoption on international as well as domestic sales. Second, our study is one of the first attempts to model the mechanisms through which AI adoption can boost firm sales performance including domestic and foreign margins, and finds that AI is not a scale-neutral

technology where its international commercial returns accrue almost exclusively to large firms with substantial intangible capital. Third, we provide causal evidence on the impact of AI on firm-level domestic and foreign sales addressing the interplay between AI adoption and sales across firm-size distribution.

Our findings have important policy implications for debates on digitalization, industrial policy, and international competitiveness. The results suggest that policies promoting AI diffusion may have uneven effects across firms, with international gains concentrated among larger firms with stronger intangible capabilities. Absent complementary policies that ease financial constraints or support intangible investment, AI adoption alone is unlikely to broaden export participation or international market access among smaller firms. From a trade policy perspective, this highlights the importance of considering how digital technologies interact with firm size and intangible assets when designing policies intended to promote international competitiveness and inclusive globalization.

Related Literature Our paper relates to and contributes to several strands of the literature. First, the literature on technology adoption and international trade has established a strong link between technological advancements and enhanced trade performance. Classic work by Grossman and Helpman (1991) set the foundation by exploring how technological progress influences trade patterns and growth through innovation and knowledge spillovers. Similarly, Eaton and Kortum (2002) developed a quantitative model showing that technology adoption and diffusion are key drivers of international trade and productivity differences across countries. Their model emphasizes how technology helps firms lower costs and improve quality, thereby enhancing export competitiveness. Further studies have examined the conditions under which firms benefit from technology in global markets. Melitz (2003) introduced a seminal model of heterogeneous firms in international trade, demonstrating how productivity gains from technology adoption lead to self-selection into exporting, with more productive firms succeeding in competitive international markets. Building on this, Lileeva and Trefler (2010) analyzed Canadian firms and found that technology investments were associated with increased productivity and export propensities, especially in firms exposed to tariff reductions. Bloom et al. (2016)

extended this understanding by examining how trade-induced competition from China prompted European firms to innovate and adopt advanced technologies, highlighting how competitive pressures drive technological upgrading as a means of survival in global markets. These foundational studies establish that technology adoption plays a critical role in boosting productivity and enhancing firms' international market access. However, most of this literature focuses on traditional technologies, with limited attention to AI specifically. Our study diverges by focusing on AI adoption's impact on firm-level domestic and international sales, contributing new insights into how AI, as a transformative technology, affects trade outcomes in ways that go beyond the effects of traditional technological advancements.

Another related strand focuses on the role of AI on firm performance. The empirical literature has largely focused on firm growth and innovation aspects. Agrawal et al. (2019) discuss the predictive capabilities of AI technologies and how it impacts firm decision-making process, concluding that AI technologies enable firms to adapt faster to changing business environment by creating new business opportunities. Cockburn et al. (2019) argue that AI technologies can substantially revolutionize the innovation and R&D processes by accelerating the accumulation of knowledge. Rock (2021) utilizes the open source launch of Google's TensorFlow as a natural experiment to show that adoption of AI technologies increases market value, while there is no evidence of significant impact on revenue productivity. Hirvonen et al. (2022) provide new evidence from Finland on the impact of adoption of AI technologies on employment, skill demand, and firm performance. Babina et al. (2024) investigate how AI investments affect firm growth and innovation outcomes, specifically examining the relationship between firms' AI investments and their subsequent growth in sales, employment, and product offerings. Goldfarb and Trefler (2018) study key features of AI with respect to roles of scale, competition, and knowledge creation and knowledge diffusion, concluding that there is still much to learn for a holistic understanding of trade effects of AI.

Lastly, we contribute to the recent research which has begun to investigate the intersection of AI, digital technologies, automation, and international trade outcomes across both macro and firm-level settings. Chander (2021) analyses AI regulation from the per-

spective of international trade law. His work suggests that WTO agreements promote international trade powered by AI. Ferencz et al. (2022) explore the interlinkages between AI technologies and international trade and outlines key trade policy considerations for policy makers seeking to harness the full potential of AI technologies. Brynjolfsson et al. (2019) study a key application of AI: machine translation, and find that the introduction of a new machine translation system has significantly increased international trade on this platform, increasing exports by 10.9%. Lo and Lee (2024) study the contribution of digitalization and AI intensities to facilitating international trade, and find for every 1% increase in a country's AI intensity that its exports rise by 1.01% to 1.30%. Xu and Tian (2023), analyse if AI improves export product quality, and they find that emerging countries should motivate their exporters to adopt intelligent manufacturing methods to enhance product quality. Jakubik et al. (2025) develop a theoretical framework in which an AI shock to production influences bilateral trade patterns and find empirical evidence that countries with higher AI exposure exhibit larger export flows, particularly among advanced economies, highlighting the role of AI in shaping global comparative advantage. Mishra et al. (2023) provide a data-driven framework that integrates AI-based specialization with goods and services export specialization to shed light on how AI adoption may shape pathways of economic diversification and comparative advantage at the national level. Barbosa (2024) documents, using Portuguese firm data, that AI and related digital technology adoption are linked to heterogeneous export performance; firms with advanced digital capabilities and higher pre-existing performance are more likely to adopt these technologies and to participate in export markets, emphasizing how pre-adoption firm characteristics shape outcomes. Similarly, Antoniades et al. (2025) combine customs data with online job postings to show that AI hiring in sales and analytics sectors mitigates information frictions and is associated with increased export scope and entry into new markets, particularly for larger and more differentiated firms. Carballo et al. (2022) show that online business platforms can reduce search and matching frictions in international trade, improving firms' export performance and generating heterogeneous effects across firms depending on their initial capabilities. Artuc et al. (2023) document that automation technologies, proxied by robot adoption, affect North–South trade patterns by

shifting both imports and exports, highlighting how technology adoption can reallocate activity across international margins.

While these contributions highlight the importance of AI for trade and highlight heterogeneity in technology adoption and export performance, they leave open how improvements in the effectiveness of AI-skilled labor translate into distinct sales margins at the firm level, especially when distinguishing foreign sales from export sales. Building on this literature, our paper exploits a plausibly exogenous improvement in AI-worker effectiveness through the 2015 release of Google TensorFlow as a shock to AI implementation costs and productivity. This allows us to identify causal effects of enhanced AI effectiveness on multiple sales components, showing that AI-driven gains in foreign sales are concentrated among large, intangible-capital-intensive firms, while systematic effects on export sales are limited. In doing so, we highlight the role of internal organizational capacity and scale in shaping how AI adoption influences international commercialization, complementing and extending existing work on AI and trade.

Layout The remainder of the paper is organized as follows. Section 2 describes the data sources, variable construction, and sample features. Section 3 presents descriptive analysis and stylized facts. Section 4 outlines the empirical strategy and identification approach. Section 5 reports the main empirical results. Section 6 presents a range of robustness checks. Finally, Section 7 concludes.

2 Data, Measurement, and Sample Composition

2.1 Data and Variable Measurement

The empirical analysis relies on three primary data sources. First, Compustat provides annual firm-level information including total sales, assets, capital expenditures, R&D spending, cash holdings, and other financial characteristics.

For firm-level domestic and foreign sales, unfortunately we do not have access to firm-to-firm trade relationships or transaction-level customs data, which prevents us from directly tracing buyer–supplier linkages or mapping firms' exposure through global value

chains. To address this limitation, we use the Compustat Historical Segment database to construct measures of firm-level foreign and domestic sales. This database reports sales by geographic segment, distinguishing between domestic and nondomestic operations. In addition, for domestic segments, Compustat separately reports export sales. Following Amihud et al. (2015), we compute firm-level foreign sales by summing sales from nondomestic segments (identified by the Compustat variable $GEOTP = 3$) and adding export sales from domestic segments. Exports are measured using the Compustat variable $SALEXG$, which captures revenue generated from goods and/or services produced domestically and sold to customers outside the firm's home country. Firm-level domestic sales are then calculated as sales from domestic segments (identified by $GEOTP = 2$) net of export sales ($SALEXG$).

Third, we use measures of AI-related workers by Babina et al. (2024), which comes from detailed individual job histories and job posting information by classifying employees according to whether their roles involve machine learning, data science, or other AI-related tasks.[1] We focus on the "broad AI" measure that captures directly technical AI engagement. The share of AI workers in 2014 serves as our primary variable measuring a firm's ex-ante AI intensity. While our data end in 2018, prior to the recent acceleration in AI adoption associated with the diffusion of large language models such as ChatGPT, this limitation is unlikely to materially affect our results. According to the "AI-skilled employee" definition used in this paper, workers who merely use AI-enabled tools without performing core AI development or implementation tasks are not classified as AI-skilled. As a result, the measure primarily captures firms' underlying technical capacity to develop and deploy AI, rather than the more recent, widespread use of off-the-shelf AI applications.

We categorize firms as "large" or "small" depending on whether their total assets fall above or below the industry-year median. Control variables include lagged total assets, the ratio of R&D expenditures to sales, the logarithm of capital expenditures, the cash-to-asset ratio, firm age, markups, and Tobin's Q. We measure firm-level markup based on the methodology of De Loecker et al. (2020). We proxy firm-level productivity using sales per

[1] For further details on the datasets and sample construction, please see Babina et al. (2024).

employee. We also construct a financial constraint indicator following the methodology of Hadlock and Pierce (2010), and intangible capital using the Peters and Taylor (2017) measures.

2.2 Sample Composition

Table A.1 presents the variable definitions, units, and the corresponding data sources in our analysis. We construct a firm-year panel covering the period 2010–2019, which provides a sufficiently long pre- and post-treatment window around the 2015 TensorFlow release. Table A.2 summarizes the sample construction procedure, and the final estimation sample consists of 3,792 firms and 27,855 firm-year observations. All monetary variables are deflated to constant 1990 U.S. dollars using the Consumer Price Index (CPI). Continuous variables are winsorized at the 1st and 99th percentiles within each year.

Table A.3 reports the sample composition by year and industry. Panel A shows that observations are relatively evenly distributed over the 2010–2019 period, with no sharp changes around the 2015 TensorFlow release, alleviating concerns about treatment-year sample selection. Panel B indicates broad industry coverage across Fama–French sectors, with substantial representation in High Tech, Healthcare, Manufacturing, and Wholesale and Retail.

Table A.4 presents summary statistics for the main variables. The distribution of sales outcomes highlights substantial firm heterogeneity and firm size, intangible capital, and R&D exhibit wide dispersion, and the share of AI workers is small on average but highly skewed. This skewness highlights the relevance of exploiting pre-treatment AI exposure in the empirical design. Table A.5 reports pairwise correlations among the variables. AI exposure is positively correlated with firm size, productivity, R&D intensity, and intangible capital, consistent with the idea that AI adoption is more prevalent among larger and more knowledge-intensive firms. At the same time, the correlations are well below unity, suggesting that AI exposure captures a distinct dimension of firm capabilities rather than merely proxying for size or innovative activity. This mitigates concerns that our results simply reflect broader firm characteristics.

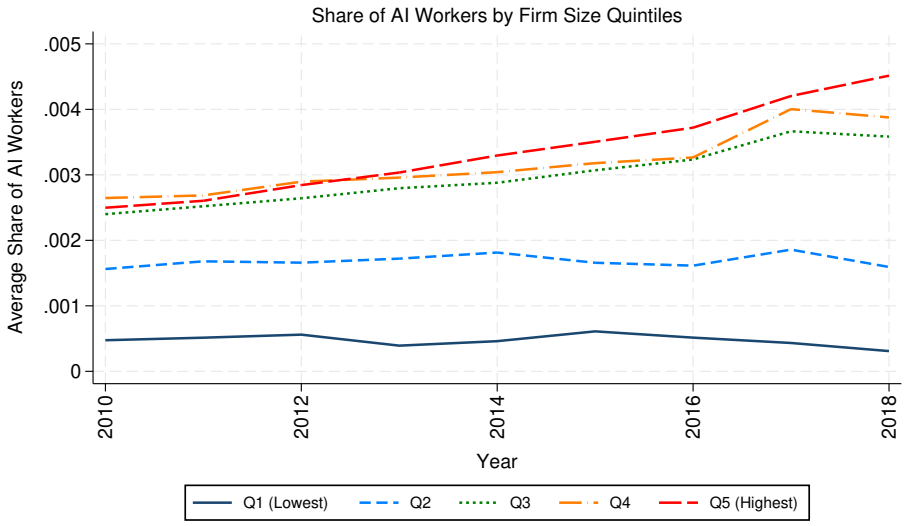
Figure B.1 presents the distribution of firm-level sales components. The histograms reveal substantial heterogeneity across firms, particularly in foreign and export sales, where a large mass of firms records zero or very low values. This skewness indicates that international activity is concentrated among a subset of firms, reinforcing the importance of accounting for firm-level heterogeneity when analyzing the impact of AI adoption on sales outcomes. Figure B.2 shows average sales components across Fama–French industries. The figure highlights pronounced cross-industry variation in the composition of sales, with some industries exhibiting greater reliance on foreign or export sales than others. This heterogeneity highlights the need to control for industry-specific factors and motivates the inclusion of industry fixed effects in the empirical analysis.

Figure B.3 illustrates the distribution of the firm-level share of AI workers. The distribution is highly skewed, with most firms employing no or very few AI-related workers and a small number of firms exhibiting relatively high AI intensity. This pattern suggests meaningful variation in firms' pre-treatment exposure to AI, which we exploit to identify heterogeneous responses to the TensorFlow shock. Figure B.4 reports the average share of AI workers across industries. The figure shows substantial differences in AI intensity across Fama–French sectors, indicating that AI adoption is not uniform across industries.

3 Descriptive Analysis

Figure 1 depicts the evolution of the share of AI workers across firm size quintiles. The share of AI workers increases steadily over time, but the growth is highly skewed toward larger firms. Firms in the top size quintiles (Q4 and particularly Q5) consistently employ a substantially higher and rising share of AI workers, while smaller firms, in particular those in Q1, show minimal adoption and even a slight decline. This pattern indicates that AI-related human capital is concentrated in larger firms, suggesting a greater capacity to invest in and integrate AI technologies compared to smaller firms.

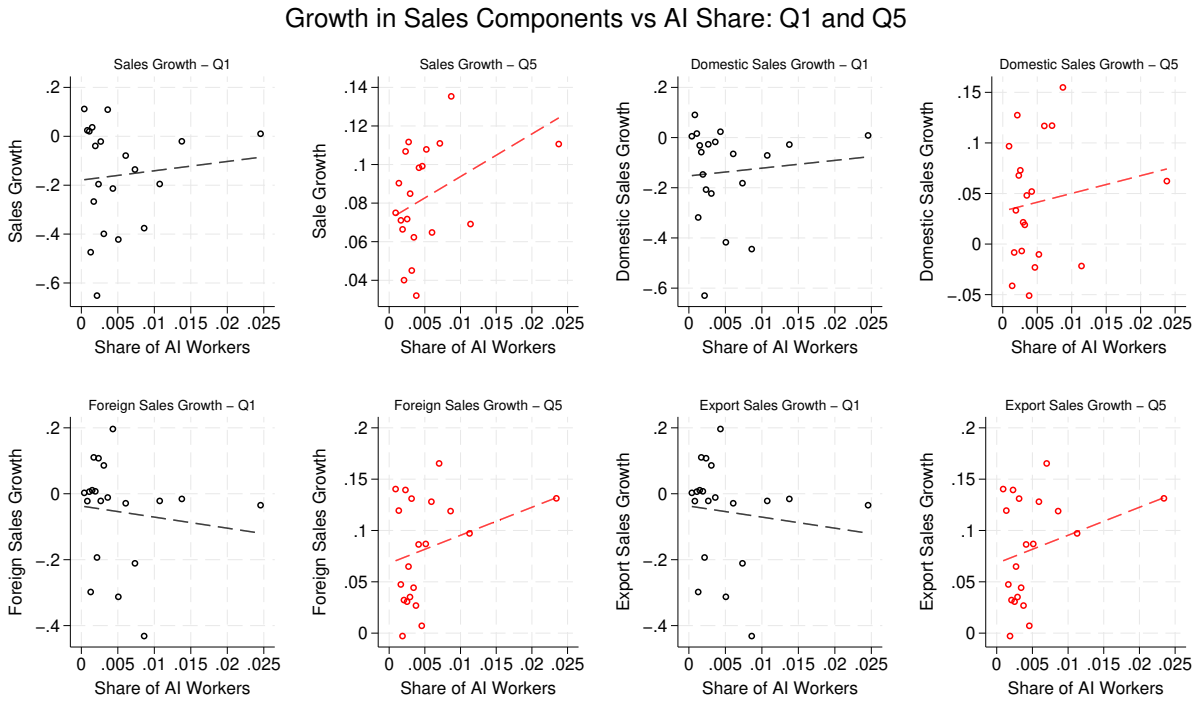
Figure 1: Share of AI Workers by Firm Size Quintiles, Over Time



Note: This figure depicts the average share of AI workers across firm size quintiles over time. Firm size quintiles are constructed based on total assets within each industry and year.

Figure 2 presents a comparison of the relationship between the share of AI workers and the growth of sales components at the two ends of the firm-size spectrum—lowest-quintile firms (Q1) and highest-quintile firms (Q5). For Q1 firms, the association between AI intensity and total, domestic, foreign, or export sales growth is weak or slightly negative. In contrast, Q5 firms exhibit a clear positive relationship across all sales components, with higher AI shares associated with faster growth. These patterns suggest that larger firms are better positioned to translate AI-related human capital into commercial performance, whereas smaller firms show negligible growth response.

Figure 2: Growth in Sales Components vs. Share of AI Workers



Note: This figure illustrates the relationship between growth in different sales components and the firm-level share of AI workers for firms in size Quintiles 1 and 5. Firm size quintiles are constructed based on total assets within each industry and year.

To examine these relationships more systematically, we estimate firm-level panel regressions of the following form:

$$\begin{aligned} Y_{i,t} =& \beta_1 \text{Share of AI Workers}_{i,t-1} + \beta_2 \left(\text{Share of AI Workers}_{i,t-1} \times \text{Log Total Assets}_{i,t-1} \right) \\ &+ X'_{i,t-1}\gamma + \lambda_{j,t} + \phi_r + \mu_i + \varepsilon_{i,t} \end{aligned} \tag{1}$$

where $Y_{i,t}$ denotes the logarithm of firm-level sales outcomes for firm i in year t, including total sales, domestic sales, foreign sales, or export sales. The main explanatory variable, $\text{Share of AI Workers}_{i,t-1}$, measures the lagged share of AI-related workers and captures firms' AI intensity prior to realized sales outcomes. The interaction term with lagged firm assets allows the association to vary with firm size.

The vector $X_{i,t-1}$ includes lagged firm-level control variables, namely the logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital

expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markups. We include firm fixed effects μ_i, industry-by-year fixed effects $\lambda_{j,t}$, and state fixed effects ϕ_r, which absorb time-invariant firm characteristics, industry-specific shocks, and time-invariant regional heterogeneity. Standard errors are clustered at the firm level.

Table 1: AI Workers and Sales Components

	(1)	(2)	(3)	(4)	(5)
Panel A: Log Sales					
L.Share of AI Workers	107.9***	26.31***	112.0***	1.706**	-2.267
	(7.937)	(2.221)	(9.041)	(0.737)	(1.833)
L.Share of AI Workers × L.Log Total Assets					0.733**
					(0.308)
Adjusted R^2	0.0353	0.907	0.139	0.978	0.978
Observations	48,454	36,305	48,445	29,830	29,830
Panel B: Log Domestic Sales					
L.Share of AI Workers	108.0***	43.56***	110.8***	1.769	-4.634
	(10.52)	(7.637)	(11.70)	(1.892)	(4.644)
L.Share of AI Workers × L.Log Total Assets					1.181
					(0.977)
Adjusted R^2	0.0316	0.446	0.0887	0.907	0.908
Observations	48,458	36,294	48,449	29,823	29,823
Panel C: Log Foreign Sales					
L.Share of AI Workers	117.8***	23.83***	104.7***	0.962	-10.09***
	(8.000)	(4.203)	(8.014)	(1.227)	(3.118)
L.Share of AI Workers × L.Log Total Assets					2.038***
					(0.638)
Adjusted R^2	0.0407	0.421	0.112	0.949	0.949
Observations	48,461	36,220	48,452	29,819	29,819
Panel D: Log Export Sales					
L.Share of AI Workers	-0.441***	-0.372*	-0.393**	0.0394	-0.528
	(0.167)	(0.211)	(0.173)	(0.308)	(0.785)
L.Share of AI Workers × L.Log Total Assets					0.103
					(0.102)
Adjusted R^2	0.000122	0.00700	0.00769	0.487	0.487
Observations	48,463	35,859	48,454	29,481	29,481
Controls	No	Yes	No	Yes	Yes
Industry-Year FE	No	No	Yes	Yes	Yes
State FE	No	No	No	Yes	Yes
Firm FE	No	No	No	Yes	Yes

Note: This table reports panel regression estimates of the relationship between the variety of sales components and the lagged share of AI workers. Control variables include lagged logarithm of assets, logarithm of R&D expenditures per sales, logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, age, and markup. Standard errors (in parentheses) are clustered at the firm level. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 1 presents the results and shows that positive association between AI exposure and firm sales is driven entirely by large firms. Across Panels A–C, the interaction between lagged AI intensity and firm size is positive, statistically significant, and economically meaningful, indicating that only firms with above-median assets convert AI capabilities into higher total, domestic, and international sales. The association is especially

pronounced for international sales, where AI–size complementarity is strongest. In contrast, small firms exhibit coefficients that are close to zero and statistically insignificant throughout, demonstrating that AI exposure does not associate with measurable sales gains for firms lacking sufficient scale. Panel D further shows that export sales do not display consistent association, suggesting that the improvements for large firms reflect broader international expansion rather than export-specific channels. Overall, Table 1 highlights a clear size-dependent pattern: only large firms are more likely to leverage AI adoption into commercial gains across both domestic and international markets.

These baseline results should be interpreted with caution due to potential endogeneity. We therefore turn to a causal inference framework in the next section.

4 Empirical Framework

4.1 Identification Strategy

To assess how firm-level AI adoption influences different components of sales performance, we build on the insight of Rock (2021) and exploit the 2015 release of Google TensorFlow as a plausibly exogenous shock to the effectiveness of AI-related human capital. TensorFlow is a widely used open-source machine-learning framework that substantially reduced the technical complexity, fixed costs, and implementation frictions associated with deploying advanced AI models. By standardizing model development, simplifying deployment, and facilitating faster experimentation, the release of TensorFlow increased the marginal productivity of existing AI-skilled workers, allowing firms to leverage their AI workforce more effectively without requiring immediate changes in firm size, capital structure, or workforce composition. In addition to Rock (2021), recent studies explicitly emphasize the role of TensorFlow as a pivotal shock in the diffusion of artificial intelligence. Using patent-level evidence, Chen et al. (2024) show that the value premium of AI innovations increases following the open-sourcing of TensorFlow, particularly in application sectors where shared machine-learning infrastructure facilitates knowledge spillovers. Complementary evidence from Chen et al. (2025) documents a sharp post-2015 increase in firms' demand for AI skills using job postings from the Lightcast database by

showing that TensorFlow-related skills are virtually absent in the pre-period (only one mention between 2010 and 2015) but surge to 26,546 mentions between 2016 and 2021. Exploiting the release of TensorFlow in a difference-in-differences framework, they provide causal evidence that AI adoption improves climate-related innovation, environmental management practices, and disclosure quality.

Importantly, the release of TensorFlow was not targeted at specific firms, industries, or regions. Instead, it constituted a broad technological improvement that enhanced the usability and productivity of pre-existing AI expertise economy-wide. As a result, firms with a higher pre-2015 share of AI-skilled workers were differentially exposed to this shock, providing a natural source of variation to identify how improvements in AI worker effectiveness translate into changes in firm-level sales outcomes.

In our setting, this shock is expected to increase the marginal effectiveness of AI-skilled workers by enabling faster model development, simpler implementation, and more efficient experimentation. Because TensorFlow primarily enhances the productivity of pre-existing AI talent, rather than inducing immediate changes in firm size, capital structure, or workforce composition, firms with a higher pre-2015 share of AI-skilled workers are plausibly more exposed to the shock. This feature motivates our use of pre-treatment AI worker intensity as a firm-level measure of exposure to the TensorFlow release.

Crucially, we do not interpret the TensorFlow shock as generating uniform or mechanical gains across firms. If the shock operated solely through generic task automation, one would expect relatively homogeneous effects. Instead, improvements in AI effectiveness are likely to translate into economic outcomes only insofar as firms are able to integrate AI technologies into their existing organizational and production structures. This reasoning motivates our emphasis on heterogeneity by firm size, particularly when examining domestic versus foreign sales responses.

A remaining concern is that firms with higher pre-2015 AI worker shares may differ systematically from other firms along observable dimensions, raising the possibility of selection into AI-related employment. While this concern does not challenge the interpre-

tation of TensorFlow as an economy-wide improvement in AI effectiveness, it motivates additional empirical safeguards. To address potential selection on observables, we complement our baseline difference-in-differences analysis with propensity score matching in the robustness section, comparing firms with similar pre-treatment characteristics but differing levels of AI exposure.

4.2 Identification Assumptions

Our identification strategy treats the 2015 release of Google TensorFlow as a plausibly exogenous technological shock that differentially affected firms based on their pre-existing AI intensity. We assume that, in the absence of TensorFlow, firms with higher and lower pre-2015 shares of AI workers would have followed parallel trends in outcomes. Under this assumption, any post-2015 divergence can be attributed to the improved effectiveness of AI adoption enabled by TensorFlow. We further assume that the share of AI workers measured in 2014 is predetermined and not influenced by anticipation of TensorFlow's release, which is plausible given the unexpected nature of its open-source launch. Finally, we assume that no other contemporaneous shocks in 2015 differentially affected firms by AI intensity. To address this concern, we include rich firm-level controls as well as firm and industry-by-year fixed effects. Together, these assumptions support the interpretation that our estimates capture the causal impact of TensorFlow through enhanced AI effectiveness, rather than through unrelated channels.

4.3 Empirical Design

We proxy firms' exposure to this exogenous shock using the pre-treatment share of AI-related workers in 2014, which serves as a measure of ex-ante AI intensity. We implement the following empirical model:

$$Y_{i,t} = \beta \left(\text{Post}_t \times \text{Share of AI Workers}_{i,2014} \right) + X'_{i,t-1}\gamma + \lambda_{j,t} + \phi_r + \mu_i + \varepsilon_{i,t}, \tag{2}$$

with the dependent variable $Y_{i,t}$ that corresponds to the firm-level sales outcomes, including the logarithm of total sales, domestic sales, foreign sales, or export sales. The indicator

variable Post_t equals one for years after the 2015 release of Google TensorFlow and zero otherwise. The variable Share of AI $\text{Workers}_{i,2014}$ measures firms' ex-ante AI exposure, defined as the share of AI-related workers prior to the TensorFlow release. The interaction term Post$t$ $\times$ Share of AI $\text{Workers}_{i,2014}$ captures the differential post-2015 change in sales outcomes for firms with greater AI readiness.

The vector $X_{i,t-1}$ includes the same set of lagged firm-level control variables such as the one-year lagged values of the logarithm of sales, the logarithm of assets, logarithm of R&D expenditures, logarithm of capital expenditures, the cash-to-asset ratio, and Tobin's Q. We include firm fixed effects μ_i, industry-by-year fixed effects $\lambda_{j,t}$, and state fixed effects ϕ_r, which together absorb time-invariant firm heterogeneity, common industry shocks, and time-invariant regional factors. Standard errors are clustered at the firm level. This specification identifies the average post-2015 effect of AI exposure on firm sales outcomes.

We also implement a dynamic difference-in-differences framework to trace the evolution of treatment effects over time through the following specification:

$$Y_{it} = \sum_{\tau \neq -1} \beta_\tau \left(\mathbb{1}\{t = \tau\} \times \text{Share of AI Workers}_{i,2014} \right) + X'_{it-1}\gamma + \lambda_{st} + \phi_r + \mu_i + \varepsilon_{it} \qquad (3)$$

with Y_{it} that denotes the outcome of firm i in year t, measured as the logarithm of either total sales, domestic sales, foreign sales, or export sales. The interaction terms $\mathbb{1}\{t = \tau\} \times$ Share of AI $\text{Workers}_{i,2014}$ represent event-time indicators, where $\tau = 0$ corresponds to the year of the TensorFlow release (2015), and $\tau = -1$ is omitted and serves as the baseline period. The other variables are same as in the regression specification (2).

5 Results

5.1 Baseline DiD

Table 2 presents baseline DiD estimates linking firm-level AI exposure to sales performance following the 2015 release of TensorFlow. Across Panels A–C, the coefficient on the interaction term Post $\times$ Share of AI Workers_{2014} is positive, precisely estimated, and eco-

nomically meaningful in all specifications, indicating that firms with greater pre-existing AI intensity experienced larger post-2015 increases in total, domestic, and foreign sales relative to less AI-intensive firms. While the estimated magnitudes decline with the inclusion of additional controls and fixed effects, the coefficients remain statistically significant once firm fixed effects, industry-year fixed effects, and state fixed effects are introduced, suggesting that the results are not driven by unobserved cross-firm or industry-level heterogeneity. Column (3), which includes firm and industry-year fixed effects, shows that AI-exposed firms experience substantially faster sales growth following the TensorFlow shock, and Column (4), which adds the full set of controls and fixed effects, confirms this result with similar magnitude and strong precision. In contrast, Panel D does not show statistically significant effects on exports, indicating that sales gains associated with AI adoption operate primarily through expansion of domestic and foreign sales rather than through export-specific channels, which may partly reflect the limited variation and lower prevalence of export activity in the sample[2]. Overall, the table provides causal evidence that the decline in AI implementation costs disproportionately benefited firms that were ex-ante more AI-ready, translating into higher post-2015 sales growth through within-firm adjustments.

To gauge the economic significance of the estimates, we focus on the full specification in Column (4) of Table 2. Because all sales variables are measured in logarithms, we interpret the estimates in percentage terms and then convert them into dollar values using the sample mean levels of sales. Using the average pre-treatment share of AI workers in 2014 (0.231 percent), the estimated coefficient implies an increase in total sales of approximately 1.7 percent following the TensorFlow release. Evaluated at the sample mean, this corresponds to an increase of roughly $2.6 million in total sales for the average firm. Disaggregating by sales components yields a similar pattern. Domestic sales increase by about 2.1 percent, which translates into an additional $0.9-1.0 million in domestic rev-

[2]Exports refer to sales of goods or services produced domestically and sold to customers located abroad. Export sales therefore reflect a firm's participation in international markets through cross-border trade from its home production base. Foreign sales, by contrast, encompass sales generated by a firm's foreign affiliates or subsidiaries operating abroad (often through foreign direct investment, FDI), as well as, exports. Foreign sales in our dataset capture revenues earned from production and sales conducted locally in foreign markets in addition to cross-border trade.

enues. Foreign sales rise by approximately 2.4 percent, corresponding to an increase of around $0.25-0.30 million at the sample mean. Although smaller in absolute dollar terms, the foreign sales effect is economically meaningful given the highly skewed distribution of international activity in the data and reflects gains concentrated among larger, internationally active firms. In contrast, we find no economically or statistically meaningful effects on export sales. Given the low prevalence and limited within-firm variation of export activity in the sample, the implied magnitude along this margin is negligible. This pattern reinforces the interpretation that the commercial gains from AI adoption operate primarily through broader domestic and foreign market expansion, rather than through export-specific channels.

Table 2: DiD Estimates: Sales and AI

	(1)	(2)	(3)	(4)
Panel A: Log Sales				
Post × Share of AI Workers$_{2014}$	16.89***	17.29***	9.661***	7.526***
	(2.995)	(3.478)	(1.430)	(1.595)
Adjusted R^2	0.937	0.936	0.978	0.978
Observations	43,913	36,986	33,099	27,864
Panel B: Log Domestic Sales				
Post × Share of AI Workers$_{2014}$	16.71***	18.45***	12.64***	9.138*
	(4.494)	(4.440)	(4.049)	(4.736)
Adjusted R^2	0.899	0.894	0.909	0.907
Observations	43,923	36,915	33,086	27,856
Panel C: Log Foreign Sales				
Post × Share of AI Workers$_{2014}$	24.08***	21.18***	12.57***	10.45***
	(2.955)	(3.740)	(3.100)	(3.498)
Adjusted R^2	0.931	0.934	0.948	0.949
Observations	43,918	37,085	33,009	27,852
Panel D: Log Export Sales				
Post × Share of AI Workers$_{2014}$	0.282	0.0930	0.638	0.313
	(0.233)	(0.321)	(0.412)	(0.475)
Adjusted R^2	0.445	0.473	0.450	0.471
Observations	43,942	36,900	32,666	27,529
Control Variables	No	No	Yes	Yes
Industry-Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on the logarithm of sales. The variable *Post* captures the period after the implementation of Google TensorFlow in 2015. The share of AI workers in 2014 is used as an ex-ante measure of firm-level AI intensity prior to this technological shift. Control variables include the one-year lagged values of the logarithm of sales, the logarithm of assets, logarithm of R&D expenditures, logarithm of capital expenditures, the cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

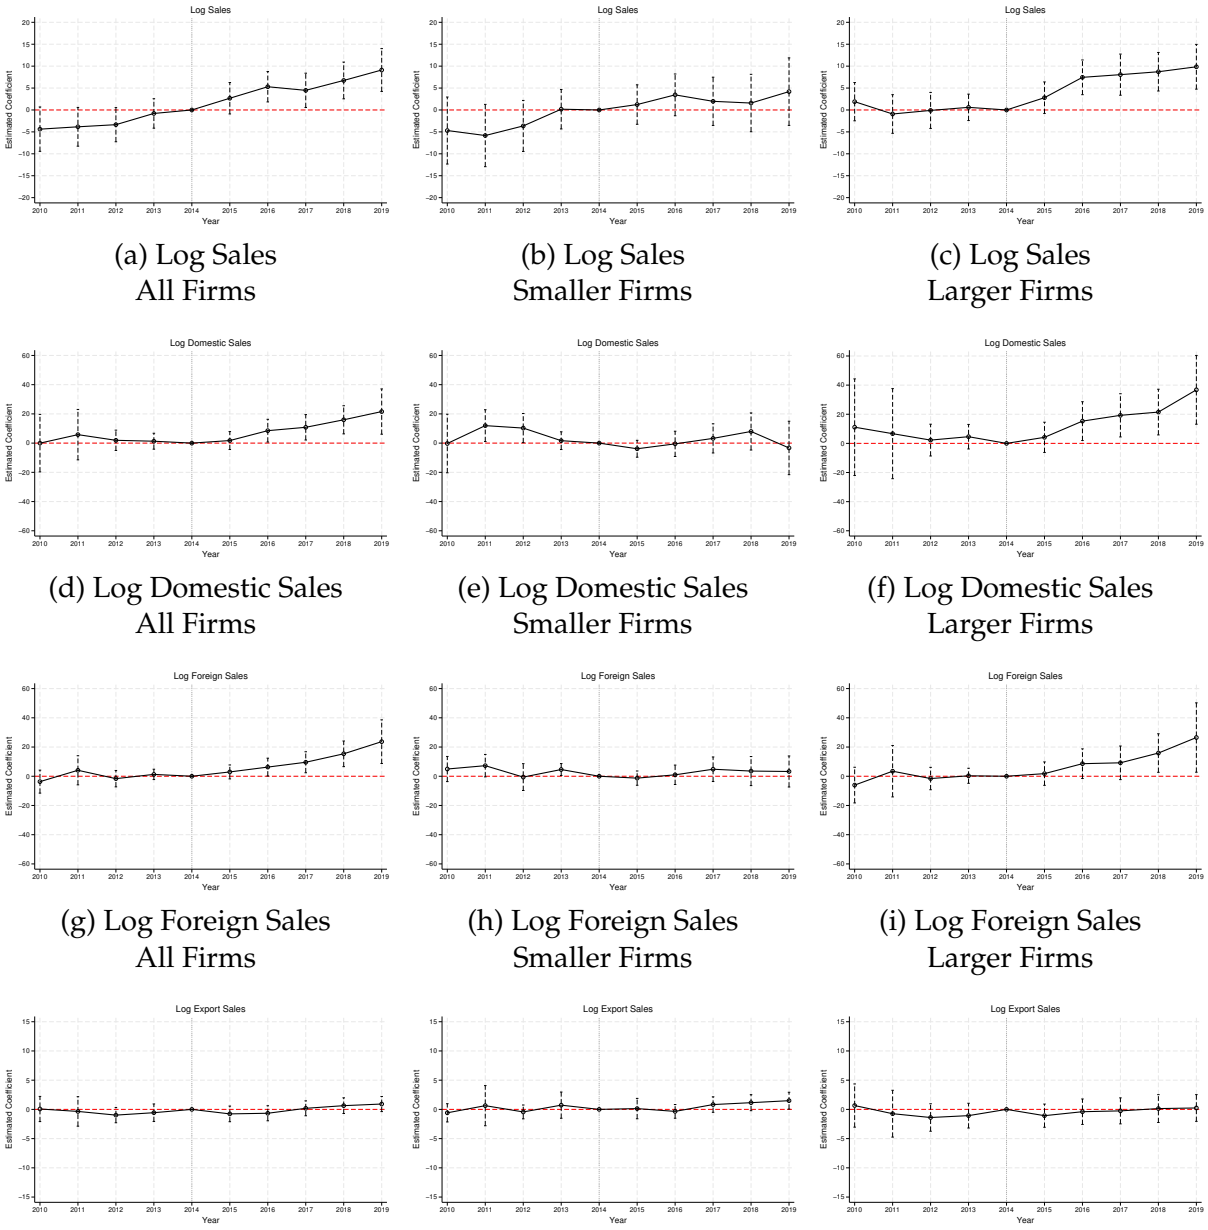
5.2 Dynamic DiD

We also conduct an dynamic difference-in-differences (DiD) analysis for our causal estimates over time to assess our parallel trend assumptions and explore the degree of persistence in the estimates. For the dynamic DiD design, we utilize a 5-year window before the 2015 release of Google TensorFlow and a 5-year window (i.e., 2015-2019, including all years) within the AI database, with 2014 serving as our base year.

Figure 3 presents dynamic DiD estimates of firms' sales responses to the TensorFlow shock, disaggregated by firm size and by sales type (total, domestic, foreign, and export). Total sales (panels (a)–(c)) rise modestly on average but increase much more strongly for larger firms, indicating that scale shapes firms' ability to translate the shock into commercial gains. Domestic sales (panels (d)–(f)) exhibit an even more pronounced divergence: larger firms experience sustained post-event growth, while smaller firms' coefficients remain close to zero, reflecting more limited adaptive capacity in home markets. Foreign sales (panels (g)–(i)) also expand only for larger firms, suggesting that international gains accrue disproportionately to firms with greater organizational and intangible resources. By contrast, export sales (panels (j)–(l)) show no systematic change for either size group, indicating that the observed increase in foreign sales does not translate into higher exports. This could be because foreign sales are driven instead by foreign production, the expansion of foreign affiliates, or the strengthening of overseas operations, rather than by physical exports themselves.

Using the 2014 mean AI-worker share (0.231%) to scale the coefficients, the implied post-treatment effect for large firms corresponds to an increase in total sales of roughly 1.6–1.9 percent, domestic sales of about 4.6 percent, and foreign sales of approximately 2.3 percent. Evaluated at the sample mean levels of sales, these effects translate into economically meaningful changes: about \$2.4–\$2.8 million in total sales, roughly \$2.0 million in domestic sales, and approximately \$0.27 million in foreign sales for larger firms. In contrast, the corresponding estimates for smaller firms are close to zero and economically negligible.

Figure 3: Dynamic DiD Estimates



(a) Log Sales All Firms

(b) Log Sales Smaller Firms

(c) Log Sales Larger Firms

(d) Log Domestic Sales All Firms

(e) Log Domestic Sales Smaller Firms

(f) Log Domestic Sales Larger Firms

(g) Log Foreign Sales All Firms

(h) Log Foreign Sales Smaller Firms

(i) Log Foreign Sales Larger Firms

(j) Log Export Sales All Firms

(k) Log Export Sales Smaller Firms

(l) Log Export Sales Larger Firms

Note: This figure presents dynamic difference-in-differences estimates for the firm-level logarithm of various sales components, stratified by firm size (measured by total assets). The below- and above-median groups are defined within industry and year. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

Overall, the dynamic DiD patterns reveal a clear asymmetry: larger firms capture the

bulk of post-event benefits, particularly in domestic and foreign markets, while smaller firms show muted or null responses. This highlights the importance of scale in absorbing technological or policy shocks and in leveraging new opportunities for market expansion, and support the baseline DiD results by showing that the average post-2015 sales effects are driven by persistent, size-dependent within-firm responses rather than transitory or pre-existing trends.

5.3 Mechanism

Table 3 presents DiD estimates of the impact of AI intensity, measured by the share of AI workers, on operating cost efficiency and labor productivity. The interaction term Post $\times$ Share of AI Workers$_{2014}$ captures whether firms with greater pre-existing exposure to AI experienced differential changes in outcomes after AI adoption relative to less AI-intensive firms. The negative coefficient in Panel A indicates that firms with higher AI intensity prior to adoption experienced larger post-adoption reductions in operating costs per sales, implying improved cost efficiency. The positive coefficients in Panel B show that these same firms also achieved larger post-adoption gains in labor productivity following AI adoption.

Table 3: DiD Estimates: Operating Cost Efficiency and Labor Productivity

Panel A: Operating Cost per Sales				
	(1)	(2)	(3)	(4)
Post $\times$ Share of AI Workers$_{2014}$	-3.814**	-4.541**	-3.094***	-2.214**
	(1.533)	(1.836)	(0.922)	(1.076)
Adjusted R^2	0.720	0.721	0.782	0.789
Observations	41,904	35,205	32,794	27,671
Panel B: Labor Productivity				
Post $\times$ Share of AI Workers$_{2014}$	8.047***	9.737***	3.931***	4.119***
	(2.235)	(2.892)	(1.266)	(1.471)
Adjusted R^2	0.899	0.902	0.953	0.956
Observations	43,584	36,581	32,889	27,707
Control Variables	No	No	Yes	Yes
Industry-Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI intensity on operating cost per sales, and labor productivity. Panel A examines the operating cost per sales; Panel B examines log labor productivity. The variable *Post* captures the period after the introduction of Google TensorFlow. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Focusing on Column (4) of Table 3, the estimates imply economically meaningful cost and productivity gains. Using the average pre-treatment share of AI workers in 2014 (0.231%), operating costs relative to sales decline by about 0.5 percentage points. Evaluated at the sample mean level of sales, this corresponds to a reduction in operating costs of approximately \$0.75-\$0.80 million per firm, or about 7 percent of average operating costs. At the same time, labor productivity increases by roughly 1 percent, which, evaluated at the mean sales level, corresponds to an increase in output on the order of \$1.4 million. These magnitudes indicate that AI adoption improves firm performance through sizable cost-efficiency and productivity channels.

Figure 4 presents dynamic DiD estimates on how firms' cost efficiency and labor productivity evolved pre and post TensorFlow shock, for all firms as well as for smaller and larger firms separately. Across all panels, pre-shock coefficients hover around zero, providing clear support for parallel trends.

The results suggest that, following the TensorFlow shock, firms with a higher pre-2014 share of AI workers experience a decline in operating costs relative to sales. The negative coefficients imply that AI-intensive firms become more cost-efficient after the shock, consistent with AI adoption lowering the marginal cost of production and improving process efficiency. This improvement in cost efficiency is driven entirely by larger firms. Smaller firms exhibit little measurable change in their cost structure, whereas larger firms show substantial reductions in operating cost per sales. This pattern suggests that larger firms are better positioned to translate early AI adoption into operational improvements, potentially due to greater complementary assets, organizational capacity, and the scale needed to absorb fixed implementation costs.

In contrast, labor productivity gradually improves after the shock, especially for larger firms. These medium-run gains are consistent with firms leveraging TensorFlow-enabled automation, analytics, or process optimization once complementary technology adoption and investments are in place.

Overall, the TensorFlow shock leads to clear improvements in firms' operating performance. AI-intensive firms experience a significant reduction in operating costs relative

to sales, indicating a marked increase in cost efficiency. These gains are mirrored on the output side: labor productivity rises substantially in the medium run, suggesting that AI adoption enhances both the efficiency of production processes and the effectiveness of the existing workforce. Importantly, both effects are driven entirely by large firms, which show pronounced cost-efficiency improvements and stronger productivity gains, whereas smaller firms exhibit little measurable response.

Figure 4: Dynamic DiD Estimates - Mechanism



(a) Operating Cost per Sales All Firms

(b) Operating Cost per Sales Smaller Firms

(c) Operating Cost per Sales Larger Firms

(d) Log Labor Productivity All Firms

(e) Log Labor Productivity Smaller Firms

(f) Log Labor Productivity Larger Firms

Note: This figure presents dynamic difference-in-differences estimates for the firm-level operating cost per sales and logarithm of labor productivity, stratified by firm size (measured by total assets). The below- and above-median groups are defined within industry and year. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

Figure 4 shows that only larger firms convert the TensorFlow shock into productivity gains, while smaller firms exhibit no meaningful adjustments. Table 4 complements this evidence by demonstrating that these gains are further amplified by intangible intensity and financial flexibility, but again only among large firms. Together, the two results highlight that the commercial payoff of AI depends critically on the scale and financial flexibility of the firms, reinforcing that TensorFlow's benefits are primarily earned by those

capable of absorbing and operationalizing the technology.

The estimates in Table 4 show that intangible-intensive firms benefit more from AI adoption, and the complementarity between AI and intangibles becomes increasingly positive with firm size, consistent with scale-dependent productivity gains. In contrast, financial constraints substantially dampen the returns to AI adoption, particularly among smaller constrained firms. However, this negative effect dissipates with size: larger constrained firms can still leverage AI effectively once scale relaxes financing frictions. Overall, the results highlight a dual mechanism, where intangible capability and financial capacity jointly determine whether AI adoption translates into sales performance.

Table 4: DiD Estimates: Sales and AI with Interaction Terms

	(1) Log Sales	(2) Log Sales	(3) Log Sales	(4) Log Sales	(5) Log Sales
Post × Share of AI Workers$_{2014}$	7.526***	7.135***	5.918	8.664***	16.95**
	(1.595)	(1.603)	(4.903)	(1.599)	(8.072)
Post × Share of AI Workers$_{2014}$ × Intangible Intensity		0.168**	-1.747**		
		(0.0856)	(0.816)		
Post × Share of AI Workers$_{2014}$ × Intangible Intensity × Size			0.817**		
			(0.364)		
Post × Share of AI Workers$_{2014}$ × Financially Constrained				-1.381	-45.76***
				(3.280)	(13.78)
Post × Share of AI Workers$_{2014}$ × Financially Constrained × Size					11.72***
					(3.687)
Control Variables	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.978	0.978	0.978	0.979	0.979
Observation	27,864	27,382	27,382	27,864	27,864

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on the logarithm of sales, interacting AI adoption with intangible intensity, financial constraints, and firm size. The variable *Post* captures the period after the implementation of Google TensorFlow in 2015. The share of AI workers in 2014 is used as an ex-ante measure of firm-level AI intensity. Control variables include lagged log sales, log assets, log R&D, log capital expenditures, cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The triple-interaction estimates reveal clear scale thresholds for the realization of AI-related gains. In Column (3), the incremental effect of intangible intensity on the AI–sales relationship becomes positive once the logarithm of firm assets exceeds approximately 2.14. Given that the sample mean of the logarithm firm assets is around 5.28, most firms lie well above this threshold, consistent with the view that intangible assets amplify AI-driven sales gains primarily among sufficiently large firms. Evaluated at the sample means and using the average 2014 AI-worker share (0.231%), the implied magnitude corresponds to an increase in sales of about 2.0 percent, or roughly $3.0 million at the mean sales level, for a one-unit increase in intangible intensity. Column (5) yields a similar im-

plication for financially constrained firms. The net effect of AI adoption turns positive once the logarithm of firm assets exceeds approximately 2.46, indicating that firm scale substantially mitigates the adverse impact of financial constraints. At the mean firm size, the implied effect is economically large for firms that are less financially constrained, with sales increasing by around 7.6 percent, corresponding to approximately $11.5 million in additional sales.

6 Robustness Checks

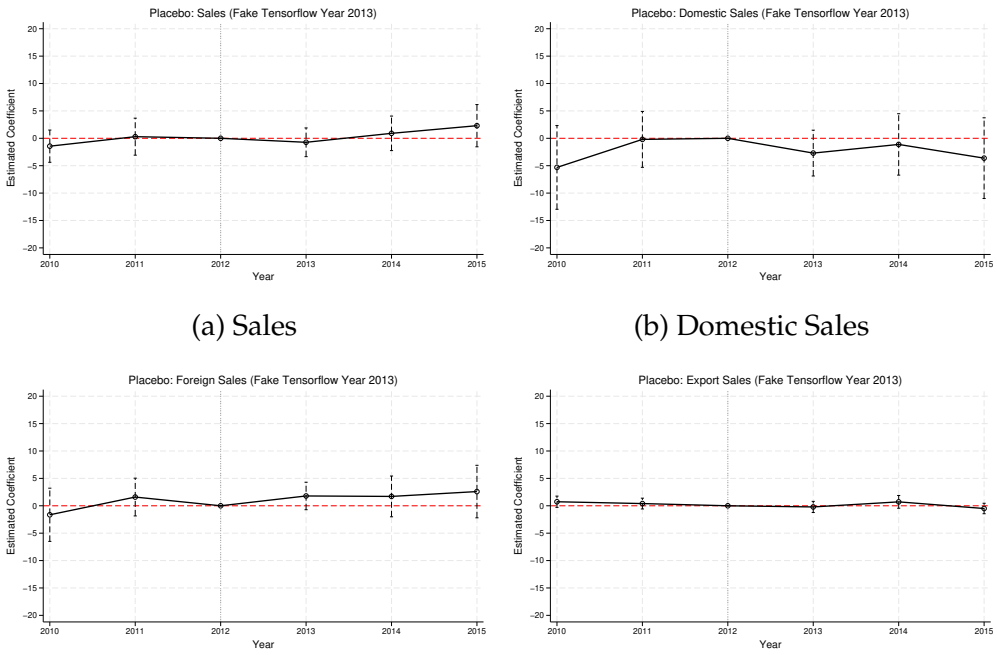
This section evaluates the stability of our baseline estimates through a set of complementary robustness checks, including placebo test, alternative treatment measures, propensity-score reweighting, additional trend specifications, and sample exclusions.

6.1 Placebo Test

To assess whether our baseline dynamic DiD estimates are driven by spurious pre-trends or shocks unrelated to the TensorFlow release, we implement a placebo test in which we artificially assign the introduction of TensorFlow to 2013 (two years before its actual release). We then re-estimate the dynamic specification using this fake treatment year. Because the actual TensorFlow release occurred in November 2015, we also include 2015 in the placebo estimation window. This ensures that no part of the real treatment period contaminates the placebo exercise and that the test cleanly isolates pre-treatment dynamics.

Figure 5 shows that the placebo coefficients fluctuate tightly around zero for all sales components—total, domestic, foreign, and export sale both before and after the placebo treatment date, including throughout the entire pre–true TensorFlow period. This confirms that firms with higher pre-2015 AI exposure did not exhibit differential sales dynamics in the years leading up to the true TensorFlow release. The absence of placebo effects strengthens the credibility of our identification strategy. It demonstrates that the post-2015 responses documented in our baseline and dynamic DiD analyses are not artifacts of pre-existing trends, anticipation effects, or correlated shocks.

Figure 5: Dynamic DiD Estimates - Placebo Test



(a) Sales

(b) Domestic Sales

(c) Foreign Sales

(d) Export Sales

Note: This figure reports dynamic difference-in-differences estimates for the firm-level logarithm of various sales components, implementing a placebo test that assigns a fake TensorFlow release year of 2013. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

6.2 Alternative AI Measure

In this section, we assess whether our findings depend on how AI exposure is measured. We replace our broad measure of AI-related workers with a narrower classification. Table 5 shows the results which remain highly consistent: firms with greater pre-2014 narrow-AI exposure experience sizable post-TensorFlow gains in total and foreign sales, with weaker or imprecise effects on domestic and export sales. This robustness confirms that our baseline findings are not driven by the specific construction of the AI exposure variable.

Table 5: DiD Estimates: Sales and Narrow AI

Panel A: Log Sales				
	(1)	(2)	(3)	(4)
Post × Share of Narrow AI Workers$_{2014}$	21.12***	28.40***	11.17***	6.979**
	(7.016)	(7.396)	(3.100)	(3.193)
Adjusted R^2	0.936	0.936	0.978	0.978
Observations	43,913	36,986	33,099	27,864
Panel B: Log Domestic Sales				
Post × Share of Narrow AI Workers$_{2014}$	17.77*	29.21**	19.33	11.50
	(10.61)	(11.46)	(12.73)	(13.73)
Adjusted R^2	0.899	0.894	0.909	0.907
Observations	43,923	36,915	33,086	27,856
Panel C: Log Foreign Sales				
Post × Share of Narrow AI Workers$_{2014}$	32.99***	30.89***	15.96*	11.47
	(5.859)	(8.874)	(8.370)	(9.075)
Adjusted R^2	0.931	0.934	0.948	0.949
Observations	43,918	37,085	33,009	27,852
Panel D: Log Export Sales				
Post × Share of Narrow AI Workers$_{2014}$	0.379*	0.342	1.102***	0.500*
	(0.195)	(0.208)	(0.341)	(0.302)
Adjusted R^2	0.445	0.473	0.450	0.471
Observations	43,942	36,900	32,666	27,529
Control Variables	No	No	Yes	Yes
Industry-Year FE	No	Yes	No	Yes
State FE	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on the logarithm of sales, using the alternative measure of narrow AI. The variable *Post* captures the period after the implementation of Google TensorFlow in 2015. The share of AI workers in 2014 is used as an ex-ante measure of firm-level AI intensity prior to this technological shift. Control variables include the one-year lagged values of the logarithm of sales, the logarithm of assets, logarithm of R&D expenditures, logarithm of capital expenditures, the cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Propensity Score Matching

To address concerns that firms with higher AI exposure differ systematically from other firms, we re-estimate the DiD after reweighting observations using propensity-score matching. We construct a propensity-score–matched sample based on firms' AI status in 2014. We identify treated firms as those that had any AI-related workers in 2014 and untreated firms as those with none. We then estimate each firm's likelihood of being treated based on observable characteristics such as R&D intensity, size (total assets), capital expenditures, cash holdings, Tobin's Q, age, and markups. These estimated probabilities serve as propensity scores and ensure treated and control firms are compared to others that looked

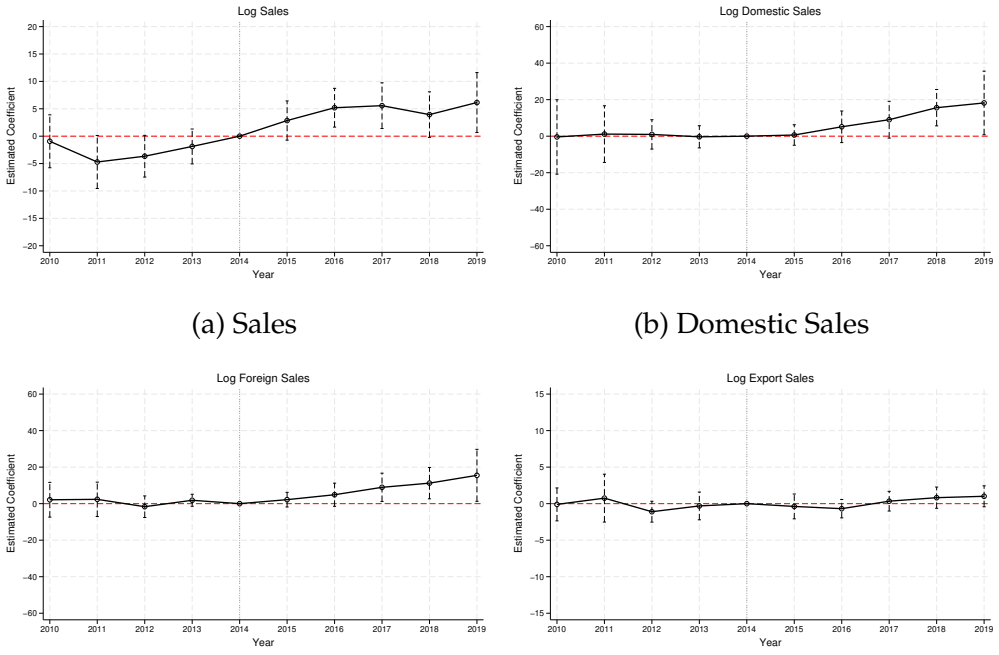
similar before the TensorFlow shock.

Table 6: DiD Estimates with Propensity Score Matching: Sales Outcomes

	(1) Log Sales	(2) Log Domestic Sales	(3) Log Foreign Sales	(4) Log Export Sales
Post × Share of AI Workers$_{2014}$	6.617***	8.472*	7.065**	0.311
	(1.795)	(4.783)	(3.509)	(0.516)
Control Variables	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Propensity Score Matching	Yes	Yes	Yes	Yes
Adjusted R^2	0.979	0.898	0.949	0.482
Observations	25,420	25,411	25,408	25,091

Note: This table reports Difference-in-Differences (DiD) estimates using the propensity score matching. Control variables include the one-year lagged values of the logarithm of sales, the logarithm of assets, logarithm of R&D expenditures, logarithm of capital expenditures, the cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 6: Dynamic DiD Estimates - Propensity Score Matching



(a) Sales

(b) Domestic Sales

(c) Foreign Sales

(d) Export Sales

Note: This figure presents dynamic DiD estimates for the logarithm of various sales components, using the propensity score matching. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

Table 6 and Figure 6 report the corresponding standard and dynamic DiD results. The key post-2015 coefficients remain similar in magnitude and significance across total, domestic, and foreign sales, while export sales again show no systematic response. These findings mitigate selection concerns and indicate that the estimated effects are not driven by differences in observable firm characteristics.

6.4 Controlling Linear Trends

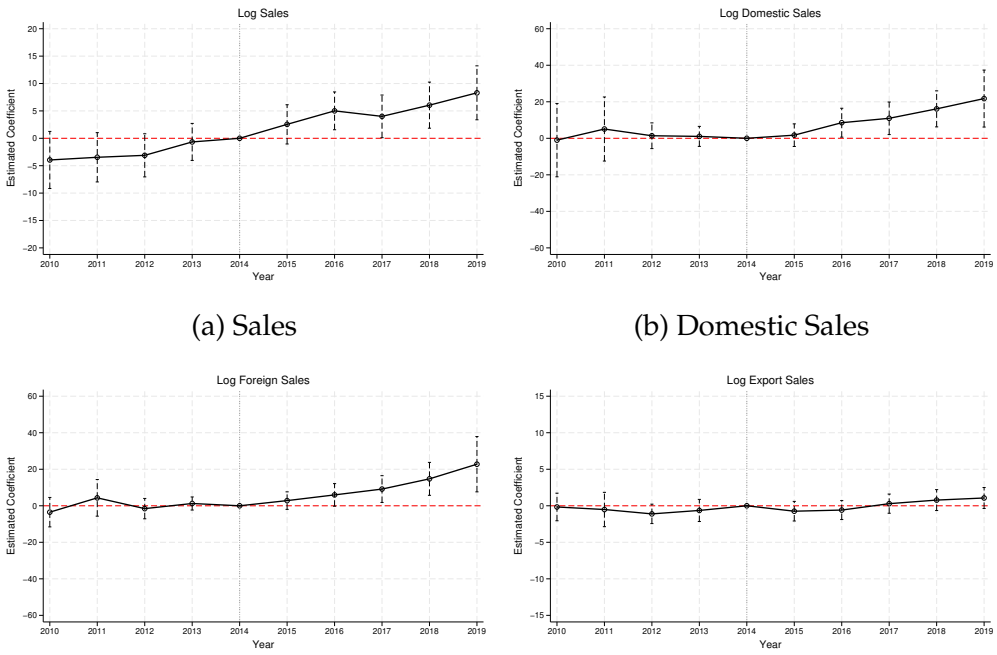
We augment the baseline specification with both state-level and year-level linear trends to ensure that differential pre-existing trajectories are not confounding the results. As shown in Table 7 for the standard DiD and Figure 7 for the dynamic DiD, the estimated post-2015 effects remain positive, stable, and statistically significant for total, domestic, and foreign sales. These patterns indicate that the results are not driven by gradual shifts correlated with geography or nationwide trends, but instead reflect the discrete change associated with the TensorFlow release.

Table 7: DiD Estimates with State and Year Trends: Sales Outcomes

	(1) Log Sales	(2) Log Domestic Sales	(3) Log Foreign Sales	(4) Log Export Sales
Post × Share of AI Workers$_{2014}$	6.750***	9.299*	9.561***	0.467
	(1.627)	(4.890)	(3.609)	(0.422)
Control Variables	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
State-level Trends	Yes	Yes	Yes	Yes
Year-level Trends	Yes	Yes	Yes	Yes
Adjusted R^2	0.978	0.907	0.950	0.474
Observations	27,864	27,856	27,852	27,529

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on sales outcomes, including both state-level and year-level trends. The variable *Post* captures the period after the implementation of Google TensorFlow in 2015. The 2014 share of AI workers is used as an ex-ante measure of firm-level AI intensity. Control variables include lagged log sales, log assets, log R&D, log capital expenditures, cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 7: Dynamic DiD Estimates - Controlling Linear Trends



(a) Sales

(b) Domestic Sales

(c) Foreign Sales

(d) Export Sales

Note: This figure presents dynamic DiD estimates for the firm-level logarithm of various sales components, including both state-level and year-level trends. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

6.5 Controlling for Other Firm-level Technologies

A potential concern is that our estimates may partly reflect broader digital transformation or correlated technological investments, rather than mechanisms specific to artificial intelligence. In particular, firms that employ a higher share of AI-skilled workers may simultaneously invest in complementary digital inputs such as information technology infrastructure, data capabilities, or data analytics that could independently affect firm performance and sales dynamics.

To address this concern, we extend our specifications by explicitly controlling for firm-level exposure to other technologies. Specifically, using the measures by Babina et al. (2024), we include controls for the shares of Information Technology (IT), data-related,

and data analytics workers, which capture general digitalization and data-driven capabilities at the firm level. These additional controls allow us to isolate the incremental effect of AI-skilled labor beyond broader digital adoption. Owing to data availability, this analysis is conducted over the 2011–2018 period.

Table 8 and Figure 8 report the corresponding baseline DiD and dynamic DiD estimates. The results remain qualitatively unchanged. Even after conditioning on other forms of technological adoption, we continue to find a positive and statistically significant post-treatment response, particularly in overall and foreign sales, for firms with higher pre-treatment AI intensity. Importantly, pre-treatment coefficients are close to zero and statistically insignificant, reinforcing the credibility of the identifying assumptions.

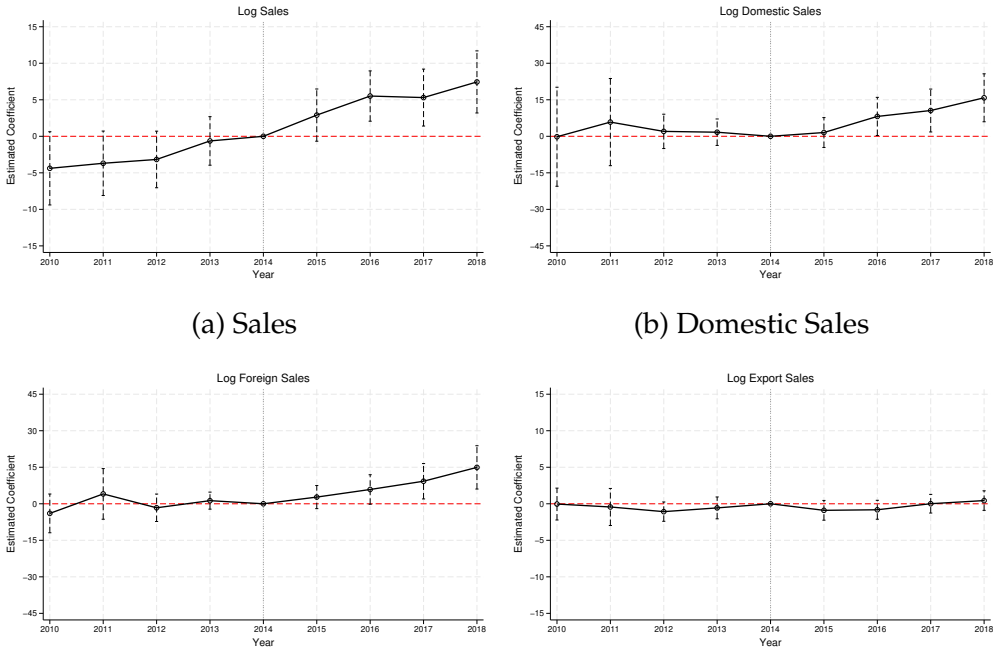
Taken together, these findings indicate that our results are not driven by generic digital transformation or contemporaneous technology investments. Instead, they point to an AI-specific channel operating through the interaction of AI-skilled labor with intangible capital and firm scale, consistent with our proposed mechanism.

Table 8: DiD Estimates by Controlling for Other Firm-level Technologies: Sales Outcomes

	(1) Log Sales	(2) Log Domestic Sales	(3) Log Foreign Sales	(4) Log Export Sales
Post × Share of AI Workers$_{2014}$	7.173***	6.743	7.658**	0.0187
	(1.575)	(5.471)	(3.637)	(0.461)
Control Variables	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
State-level Trends	Yes	Yes	Yes	Yes
Year-level Trends	Yes	Yes	Yes	Yes
Adjusted R^2	0.979	0.912	0.954	0.494
Observations	25,328	25,321	25,316	25,027

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on sales outcomes, controlling for other firm-level technologies such as share of Information Technology (IT), share of data, and share of data analytics. Since these additional controls are available until 2018 in the data, our final period in these DiDs is 2018. The variable *Post* captures the period after the release of Google TensorFlow in 2015. The 2014 share of AI workers is used as an ex-ante measure of firm-level AI intensity. Control variables include lagged log sales, log assets, log R&D, log capital expenditures, cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 8: Dynamic DiD Estimates - Controlling for Other Firm-level Technologies



(a) Sales (b) Domestic Sales

(c) Foreign Sales (d) Export Sales

Note: This figure presents dynamic DiD estimates for the firm-level logarithm of various sales components, controlling for other firm-level technologies such as share of Information Technology (IT), share of data, and share of data analytics. Since these additional controls are available until 2018 in the data, our final period in these DiDs is 2018. Control variables include the lagged logarithm of total assets, the logarithm of R&D expenditures relative to sales, the logarithm of capital expenditures, the cash-to-asset ratio, Tobin's Q, firm age, and markup. The specification includes industry-by-year, state, and firm fixed effects. Standard errors (in parentheses) are clustered at the firm level. Confidence intervals are displayed at the 95% level.

6.6 Excluding Technology Intensive Industries

Finally, we exclude high-tech and healthcare industries (sectors inherently more exposed to AI) to ensure our findings do not simply reflect sector-specific technology trends. Table 9 shows that the results remain consistent: firms with greater pre-2014 AI exposure continue to experience large and significant gains in total and foreign sales, even within this reduced sample, while domestic and export sales remain weaker or insignificant. This analysis reinforces that our conclusions are not driven by a narrow set of AI-focused industries but reflect broader firm-level heterogeneity in the ability to scale AI adoption.

Table 9: DiD Estimates without AI Intensive Industries: Sales Outcomes

	(1) Log Sales	(2) Log Domestic Sales	(3) Log Foreign Sales	(4) Log Export Sales
Post × Share of AI Workers$_{2014}$	7.375***	6.624	17.87***	0.930
	(2.798)	(10.82)	(6.470)	(1.286)
Control Variables	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Adjusted R^2	0.982	0.893	0.943	0.484
Observations	16,846	16,837	16,834	16,586

Note: This table reports Difference-in-Differences (DiD) estimates of the effect of firm-level AI adoption on sales outcomes, excluding High Tech and Healthcare industries. The variable *Post* captures the period after the release of Google TensorFlow in 2015. The 2014 share of AI workers is used as an ex-ante measure of firm-level AI intensity. Control variables include lagged log sales, log assets, log R&D, log capital expenditures, cash-to-asset ratio, and Tobin's Q. Standard errors (in parentheses) are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7 Conclusion

This paper provides new causal evidence on how artificial intelligence affects firm-level commercial performance, with a particular focus on domestic and international sales. Exploiting the 2015 release of Google TensorFlow as an exogenous reduction in the cost of AI implementation, we show that AI adoption generates substantial gains in total, domestic, and foreign sales, but only for large firms. Smaller firms, even when they employ AI-related workers, do not experience comparable improvements. This stark asymmetry highlights that AI is not a scale-neutral technology and implies that its commercial returns depend critically on firms' existing organizational capacity, intangible assets, and ability to reorganize production processes.

Mechanism analyses further demonstrate that these sales gains are underpinned by improvements in cost efficiency and labor productivity, which again accrue exclusively to large firms. Intangible capital amplifies the returns to AI adoption, while financial constraints dampen them though this negative effect is mitigated by firm size. Together, these findings suggest that AI acts as a force multiplier for firms that already possess scale, financial flexibility, and complementary intangible investments, rather than as a leveling technology that benefits firms uniformly across the size distribution.

Our results also carry important policy implications, suggesting that, without targeted

support for AI adoption and diffusion, advances in AI may further strengthen market concentration and widen productivity gaps. This concern is echoed in the World Trade Report 2025: Trade and Artificial Intelligence (World Trade Organization (2025)), which warns that early AI adopters—typically large firms and high-income economies—are likely to capture a disproportionate share of the gains in the absence of policies that account for firm-level heterogeneity. In line with the report's recommendations, our findings emphasize the need for proactive measures that broaden access to AI and support complementary investments, ensuring that the benefits of AI adoption are more widely shared across firms and economies.

References

Agrawal, A., J. S. Gans, and A. Goldfarb (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives 33*(2), 31–50.

Amihud, Y., E. Bartov, and B. Wang (2015). The pricing of corporate foreign sales risk. In *Asian Finance Association (Asian FA) 2014 Conference Paper (SSRN). Available at http://ssrn.com/abstract*, Volume 2356748.

Antoniades, A., C. He, Z. Liang, and M. Xu (2025). Learning by AI: the impact of AI hiring on Chinese exporters. *Available at SSRN 5361273*.

Artuc, E., P. Bastos, and B. Rijkers (2023). Robots, tasks, and trade. *Journal of International Economics 145*, 103828.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics 151*, 103745.

Barbosa, N. (2024). Artificial intelligence and exporting performance: Firm-level evidence from Portugal. *GEE Papers 183, Gabinete de Estratégia e Estudos, Ministério da Economia*.

Bloom, N., M. Draca, and J. Van Reenen (2016). Trade induced technical change? the impact of Chinese imports on innovation, IT and productivity. *The Review of Economic Studies 83*(1), 87–117.

Brynjolfsson, E., X. Hui, and M. Liu (2019). Does machine translation affect international trade? evidence from a large digital platform. *Management Science 65*(12), 5449–5460.

Carballo, J., M. R. Chatruc, C. S. Santa, and C. V. Martincus (2022). Online business platforms and international trade. *Journal of International Economics 137*, 103599.

Chander, A. (2021). Artificial intelligence and trade. In M. Burri (Ed.), *Big Data and Global Trade Law*, pp. 115–127. Cambridge University Press.

Chen, S.-S., Y.-S. Chen, and S.-C. Peng (2025, August). The double-edged role of Artificial Intelligence in corporate decarbonization. Working paper, SSRN. Available at SSRN: https://ssrn.com/abstract=5362708.

Chen, W. X., T. T. Shi, and S. Srinivasan (2024, May). The value of AI innovations. Working Paper 24-069, Harvard Business School.

Cockburn, I. M., R. Henderson, and S. Stern (2019). The impact of artificial intelligence on innovation: An exploratory analysis. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, pp. 115–146. University of Chicago Press.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics 135*(2), 561–644.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Ferencz, J., J. L. González, and I. O. García (2022). Artificial intelligence and international trade: Some preliminary implications. *OECD Trade Policy Papers* (260).

Goldfarb, A. and D. Trefler (2018). AI and international trade. Working Paper 24254, National Bureau of Economic Research.

Grossman, G. M. and E. Helpman (1991). Trade, knowledge spillovers, and growth. *European Economic Review 35*(2-3), 517–526.

Hadlock, C. J. and J. R. Pierce (2010). New evidence on measuring financial constraints: Moving beyond the KZ index. *Review of Financial Studies 23*(5), 1909–1940.

Hirvonen, J., A. Stenhammar, and J. Tuhkuri (2022). New evidence on the effect of technology on employment and skill demand. *Available at SSRN 4081625*.

Jakubik, A., L. Rotunno, and A. Saini (2025). Foresee the unseen: Evaluating the impact of artificial intelligence on international trade. *Journal of Policy Modeling*.

Lileeva, A. and D. Trefler (2010). Improved access to foreign markets raises plant-level productivity… for some plants. *The Quarterly Journal of Economics 125*(3), 1051–1099.

Lo, C. P. and Y. Lee (2024). Digitalization, AI intensity, and international trade. *Annals of Economics and Finance 25*(1), 251–273.

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

Mishra, S., R. Koopman, G. De Prato, A. Rao, I. Osorio-Rodarte, J. Kim, N. Spatafora, K. Strier, and A. Zaccaria (2023). AI specialization for pathways of economic diversification. *Scientific Reports 13*(1), 19475.

Peters, R. H. and L. A. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics 123*(2), 251–272.

Rock, D. (2021). Engineering value: The returns to technological talent and investments in artificial intelligence. *Available at SSRN 3427412*.

World Trade Organization (2024). Trading with intelligence: How AI shapes and is shaped by international trade.

World Trade Organization (2025). World Trade Report 2025: Making trade and AI work together to the benefit of all.

Xu, X. and C. Tian (2023). Does artificial intelligence improve the quality of export products? evidence from China. *Applied Economics Letters*, 1–5.

Appendix A Additional Tables

Table A.1: Variable Definitions

Variable	Definition (Units)	Database/Source
Sales	Sales/Turnover (Net) ($sale$) ($ million)	Compustat
Export Sales	Revenue generated from the export of goods and/or services produced domestically and sold to customers outside the company's home country ($SALEXG$) ($ million)	Compustat Historical Segment
Domestic Sales	Sales ($sales$) from domestic segments ($GEOTP = 2$) ($ million) after substracting export sales ($SALEXG$) ($ million)	Compustat Historical Segment
Foreign Sales	Sum of the sales ($sales$) from nondomestic segments ($GEOTP = 3$) ($ million) and export sales from the domestic segments ($SALEXG$) ($ million)	Compustat Historical Segment
R&D	Research and Development expense (xrd) ($ million)	Compustat
Capex	Capital expenditures ($capex$) ($ million)	Compustat
Cash-to-Asset Ratio	Cash (ch) ($ million) divided by assets (at) ($ million)	Compustat
Tobin's Q	(Assets (at) ($ million) + market value of equity ($prcc \times csho$) ($ million) – book equity (ceq) ($ million)), divided by assets (at) ($ million)	Compustat
Markup	0.85 $\times$ input revenue share (sales ($sale$) ($ million) per cost of goods sold ($cogs$) ($ million))	Compustat, De Loecker et al. (2020)
Firm Age	Difference between current year and firm's founding year	Compustat, Jay R. Ritter's Website
Labor Productivity	Sales ($sale$) ($ million) per number of employees (emp) (thousands)	Compustat
Operating Cost per Sales	Sum of cost of goods sold ($cogs$) ($ million), and selling, general and administrative Expense ($xsga$) ($ million), divided by lagged sales ($sale$) ($ million)	Compustat
Intangible Capital	Stock of intangible assets constructed using perpetual inventory method ($ million)	Peters and Taylor (2017)
Share of AI Workers	Share of number of employees with AI-related skills per total number of employees	Babina et al. (2024)

Notes: This table provides the variables used in our analysis, their definitions, units, and the corresponding data sources.

Table A.2: Sample Construction: Firm-Year Panel

Step	Dropped Firms	Dropped Observations	Remaining Firms	Remaining Observations
Initial sample (2010–2019)			15,168	93,427
Exclude Financials, Utilities, and Other Industries	6,534	41,307	8,634	52,120
After dropping singletons and missing values in regressions	4,842	24,265	3,792	27,855
Final sample			**3,792**	**27,855**

Note: This table reports the sample selection for the 2010-2019 firm-year panel used in the analysis. Industry exclusions follow standard practice: financials (SIC 6000–6999), utilities (SIC 4000–4999), and other industries with SIC $\geq$ 9000.

Table A.3: Sample Composition by Year and Industry

Panel A: Sample Composition by Year		
Year	Number of Observations	Percentage of Observations
2010	2,433	8.7
2011	2,552	9.2
2012	2,703	9.7
2013	2,916	10.5
2014	3,122	11.2
2015	3,091	11.1
2016	3,012	10.8
2017	2,851	10.2
2018	2,685	9.6
2019	2,490	8.9
Total	27,855	100.0
Panel B: Sample Composition by Industry		
Industry	Number of Observations	Percentage of Observations
Consumer Non-Durables	1,549	5.6
Consumer Durables	904	3.2
Manufacturing	3,214	11.5
Oil, Gas, and Coal	1,864	6.7
Chemicals	1,002	3.6
High Tech	6,415	23.0
Telephone and Television	751	2.7
Wholesale and Retail	3,073	11.0
Healthcare	4,598	16.5
Other	4,485	16.1
Total	27,855	100.0

Note: This table summarizes the sample distribution by year (Panel A) and by Fama–French industry classification (Panel B).

Table A.4: Summary Statistics

	Mean	SD	P25	P50	P75	Min	Max	Count
Log Sales	5.008	2.473	3.208	5.360	6.906	0.000	10.595	27,855
Log Domestic Sales	3.786	2.887	0.005	4.247	6.209	0.000	9.903	27,847
Log Foreign Sales	2.451	2.822	0.000	0.727	4.959	0.000	9.901	27,830
Log Export Sales	0.020	0.211	0.000	0.000	0.000	0.000	3.845	27,492
Log Assets	5.284	2.393	3.512	5.503	7.083	0.000	10.931	27,855
Log R&D	-3.730	2.606	-5.797	-3.571	-1.751	-9.686	3.421	27,758
Log Capex	2.337	1.946	0.434	2.124	3.777	0.000	8.026	27,836
Cash-to-Assets Ratio	0.162	0.181	0.035	0.101	0.216	0.000	0.992	27,781
Markup	1.731	1.420	1.072	1.327	1.889	0.000	15.085	27,473
Tobin's Q	6.420	32.596	1.342	2.197	4.114	-0.086	1,482.161	27,683
Firm Age	3.226	0.786	2.773	3.367	3.807	0.693	4.779	27,841
Log Labor Productivity	3.831	1.614	2.937	4.234	4.956	-0.050	7.623	27,679
Operating Cost per Sales	0.070	0.576	-0.199	-0.065	0.127	-1.255	4.377	27,638
Log Intangible Capital	4.695	2.139	3.065	4.791	6.250	0.000	10.354	27,778
Share of AI Workers	0.0028	0.0053	0.0000	0.0008	0.0035	0.0000	0.1577	27,855

Note: This table reports descriptive statistics for variables used in the analysis. All nominal variables are deflated using the CPI (1990 = 100).

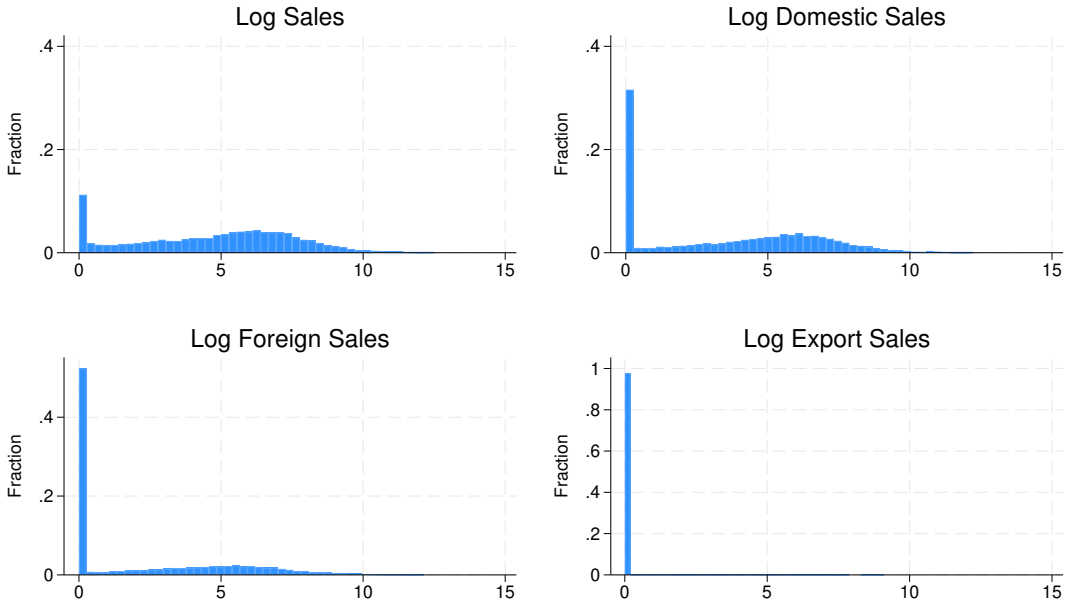
Table A.5: Pairwise Correlation Matrix

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
(1) Log Sales	1.0000														
(2) Log Domestic Sales	0.7356*	1.0000													
(3) Log Foreign Sales	0.5950*	0.5853*	1.0000												
(4) Log Export Sales	-0.0252*	-0.0022	-0.0305*	1.0000											
(5) Log Assets	0.9353*	0.6843*	0.5846*	-0.0373*	1.0000										
(6) Log R&D	-0.7632*	-0.5351*	-0.2229*	0.0086	-0.6562*	1.0000									
(7) Log Capex	0.8368*	0.6036*	0.5030*	-0.0406*	0.8823*	-0.6424*	1.0000								
(8) Cash-to-Assets Ratio	-0.3901*	-0.2902*	-0.1379*	-0.0077	-0.3429*	0.4596*	-0.3618*	1.0000							
(9) Markup	0.0163*	0.0197*	0.0743*	-0.0193*	0.0487*	0.1350*	0.0339*	0.0884*	1.0000						
(10) Tobin's Q	-0.2086*	-0.1412*	-0.0923*	-0.0121*	-0.2251*	0.1650*	-0.1319*	0.0626*	-0.0020	1.0000					
(11) Firm Age	0.4710*	0.3854*	0.3122*	0.0452*	0.4160*	-0.3593*	0.3403*	-0.1822*	-0.0550*	-0.1367*	1.0000				
(12) Log Labor Productivity	0.9002*	0.6604*	0.4943*	-0.0037	0.8449*	-0.6893*	0.7102*	-0.3674*	0.0802*	-0.2394*	0.4010*	1.0000			
(13) Operating Cost per Sales	-0.4970*	-0.3638*	-0.2800*	-0.0246*	-0.3858*	0.4944*	-0.3863*	0.3536*	-0.1718*	0.1099*	-0.3075*	-0.5211*	1.0000		
(14) Log Intangible Capital	0.8320*	0.6431*	0.6090*	-0.0475*	0.8604*	-0.4572*	0.6900*	-0.2026*	0.0792*	-0.1741*	0.4176*	0.6974*	-0.2478*	1.0000	
(15) Share of AI Workers	0.1527*	0.1281*	0.2125*	-0.0174*	0.1643*	0.1073*	0.0876*	0.0906*	0.1979*	-0.0357*	0.0487*	0.1763*	-0.0507*	0.2421*	1.0000

Note: This table reports pairwise correlations. Variables are winsorized at the 1st and 99th percentiles. An asterisk (*) indicates significance at the 10% level.

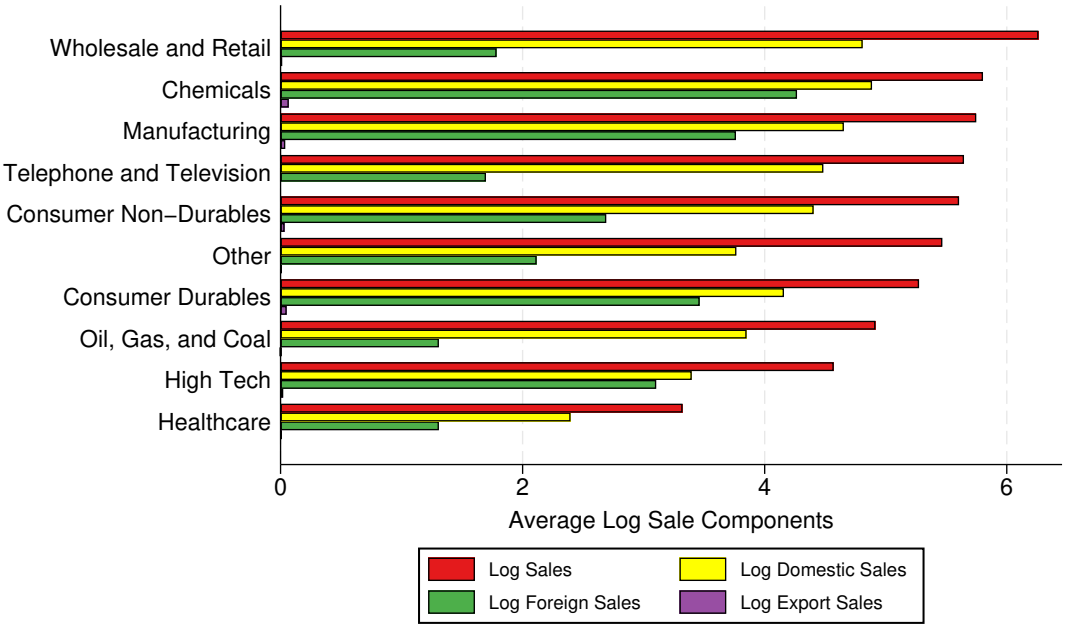
Appendix B Additional Figures

Figure B.1: Histogram of Sales Components



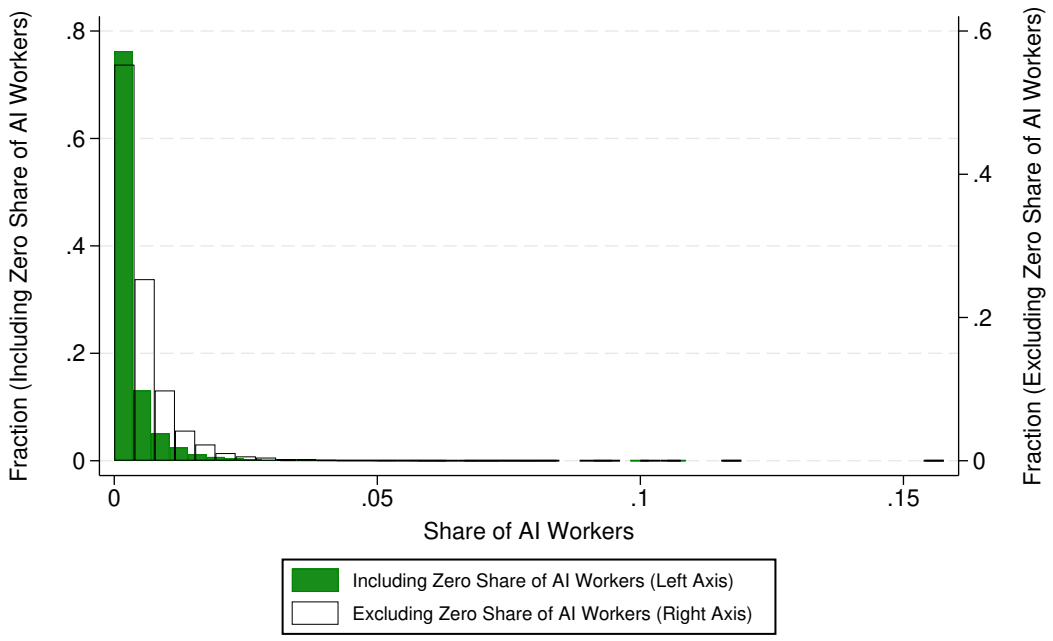
Note: This figure shows histograms of firm-level sales components.

Figure B.2: Average Sales Components by Industries



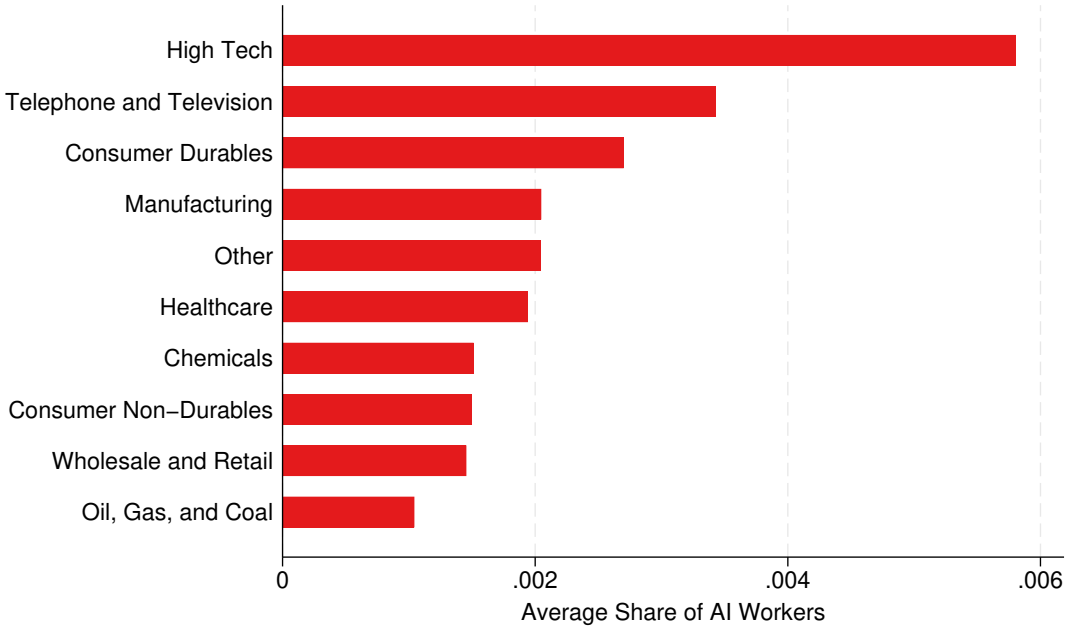
Note: This figure shows average sales components by Fama-French industries.

Figure B.3: Histogram of Share of AI Workers



Note: This figure shows histogram of firm-level share of AI workers, including and excluding the number of zero share of AI workers.

Figure B.4: Average Share of AI Workers by Industries



Note: This figure shows average share of AI workers by Fama-French industries.