Generative AI, Export Prices, and Quality-Biased Selection

Ningyuan Jia* Jia Yang* Faqin Lin† Yunfan Jia* Xirong Chen*

May 20, 2026

Abstract

This paper studies generative AI as a quality-biased shock to U.S. export prices. We incorporate the generative AI shock into a multi-country, multi-sector heterogeneous-firm model with quality sorting to study its effect on exports. The model predicts that the shock raises export unit values by selecting lower-quality varieties out of foreign markets, most strongly in AI-intensive destinations. We test this prediction with a difference-in-differences design built on the November 2022 release of ChatGPT and product exposure to generative AI. A one-standard-deviation increase in exposure raises export unit values by 2.6–3.3 percent. The effect is concentrated among initially low-price products, whose export quantities also decline, consistent with quality-biased selection. We estimate the model and find that quality-based export selection accounts for most of the unit-value response. Quantitative results imply that the global generative AI shock raised U.S. welfare by 0.53 percent.

Keywords: Generative AI, Export Selection, Export Unit Values, Quality Sorting

*University of International Business and Economics. Emails: Ningyuan Jia, jia_ningyuan@uibe.edu.cn; Jia Yang, yangjia@uibe.edu.cn; Yunfan Jia, 202400172012@uibe.edu.cn; Xirong Chen, xirongchenuibe@163.com.

†China Agricultural University. Email: linfaqin@126.com.

1 Introduction

Generative AI has quickly become one of the most important new technologies for production, competition, and international trade.[1] Two features make generative AI (GenAI) a distinctive shock for international trade. First, its diffusion was unusually rapid and global. ChatGPT reached mass scale within months. U.S. adoption has been faster than early internet adoption, and cross-country evidence documents GenAI use across more than 100 economies soon after its release (Bick, Blandin, and Deming, 2024; Ali, Mustafa, and Aysan, 2025). Second, existing evidence suggests that GenAI often raises performance most where baseline performance is lower. Experiments and workplace studies find larger gains for lower-performing or less experienced workers in writing, customer service, and knowledge tasks (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025; Dell'Acqua, McFowland III, Mollick, Lifshitz, Kellogg, Rajendran, Krayer, Candelon, and Lakhani, 2026). Field evidence from online retail also finds larger gains among smaller sellers (Fang, Yuan, Zhang, Donati, and Sarvary, 2025). After the release of ChatGPT, U.S. export unit values rose while export quantities fell (see Figure 1). This paper studies how GenAI affects U.S. export prices and identifies the channel through which these effects operate.

To answer this question, we build a quality-augmented heterogeneous-firm trade model. U.S. establishments differ in both cost and quality, so higher-quality varieties command higher unit values and can profitably serve more foreign markets. An observed export unit value is therefore a composition outcome that depends on which varieties survive selection into a destination. GenAI enters through two adoption margins. U.S. adoption raises exporter productivity, while destination-market adoption strengthens foreign competition. The shock is quality-biased in both places, with larger effects on lower-quality varieties. These two margins move export thresholds in opposite directions. U.S. adoption lowers the threshold for U.S. exporters, while destination adoption raises the threshold they face abroad. The

[1] See Goldfarb and Trefler (2019), Brynjolfsson, Li, and Raymond (2025), Jakubik, Rotunno, and Saini (2025), and Azoulay, Krieger, and Nagaraj (2024).

resulting cutoff movements change the composition of active exports and therefore the unit values recorded in customs data.

This structure yields two empirically testable implications. First, export prices and quantities depend on the relative strength of destination-side and U.S.-side GenAI adoption (Proposition 1). Destination adoption tightens the export-selection margin faced by U.S. exporters, displacing lower-quality, lower-priced varieties and raising observed export unit values. U.S. adoption works in the opposite direction by raising exporter productivity. The net price and quantity responses therefore depend on which force is stronger. The selection force is strongest for products with greater GenAI exposure, for products with low pre-release unit values, and in destinations with higher AI intensity. Second, establishment entry is governed by the same relative-shock logic (Proposition 2). U.S.-side productivity gains raise expected profits and encourage entry, while destination-side competitive pressure lowers expected export profits and discourages entry. These implications guide the empirical design.

We test these predictions with a difference-in-differences design built on the November 30, 2022 release of ChatGPT. We construct product-level GenAI exposure by mapping the industry GenAI-exposure measure of Felten, Raj, and Seamans (2021) to HS10 products, and measure destination AI intensity from pre-release OECD/LinkedIn indicators. The analysis combines monthly U.S. customs records with annual industry data.

The reduced-form estimates are consistent with the model's predictions. Products with greater GenAI exposure experience higher export unit values after the release, without a corresponding increase in export quantities. The price increase concentrates among initially low-priced products and in AI-intensive destinations, and is largest where the two coincide. Cross-destination price dispersion narrows, consistent with the exit of the lowest-priced shipments. Unlike prices, export quantities rise only in low-AI destinations, leaving the aggregate response flat. In exposed industries, the number of establishments rises while average employment per establishment falls, leaving total industry employment unchanged.

This joint pattern is difficult to reconcile with a simple productivity expansion, which would lower prices and raise quantities, or with a demand expansion, which would raise quantities. Instead, it points to quality-biased export selection.

We calibrate the model to match the pre-shock economy and estimate the AI-shock and quality-biased elasticities by simulated method of moments, matching the reduced-form price and quantity coefficients. The quantitative analysis delivers two main results. First, the decomposition shows that foreign GenAI adoption is the main force behind the export-price response. U.S. adoption lowers export unit values through its productivity effect, while foreign adoption raises them by strengthening selection against lower-quality, lower-price U.S. varieties. The quality-biased component of the shock is central to this result: when the model removes the quality gradient in AI exposure, it no longer reproduces the observed price-selection pattern. Second, the estimated model quantifies the welfare effect of the global GenAI shock. Combining the U.S. productivity gain with the foreign-selection effect, the model implies a U.S. welfare gain of 0.53 percent.

This paper contributes to three strands of literature. First, it relates to research on digital technologies and information frictions in cross-border trade. This literature treats digital technologies as reductions in the cost of searching for partners, matching, and verifying information, and finds that the internet, online platforms, and machine translation raise trade by lowering these frictions.[2] Because these technologies mainly help buyers and sellers find each other, they predict broader market access and stronger export expansion. We document a different trade margin of GenAI. Exposed products record higher export unit values without a parallel rise in quantities, and the response concentrates among initially low-unit-value products in high-AI destinations, a price–quantity wedge consistent with quality-biased export selection rather than a friction-reduction or exporter-productivity channel.

[2]Goldfarb and Tucker (2019) frame digitization as a reduction in the cost of storing, transmitting, and verifying information. On information frictions in differentiated-product trade and on the trade effects of these digital technologies, see Rauch (1999); Allen (2014); Steinwender (2018); Hjort and Poulsen (2019); Freund and Weinhold (2004); Blum and Goldfarb (2006); Lendle, Olarreaga, Schropp, and Vézina (2016); Brynjolfsson, Hui, and Liu (2019).

Second, the paper speaks to the literature on export selection, product quality, and unit values, to which it adds a destination-side mechanism. Heterogeneous-firm models show how trade shocks reallocate activity across firms, products, and destinations, with uneven adjustment across plants, workers, and regions.[3] A parallel literature reads customs unit values as signals of product quality, finding that higher-priced exporters sell more and enter richer and more difficult markets and that quality choices respond to market access and competition.[4] We add a force on the destination side. Because GenAI diffuses to destination markets as well as to U.S. exporters and improves lower-quality foreign producers more strongly, lower-quality U.S. varieties face tighter selection abroad, which reduces low-price exports and raises the average unit value of the surviving export bundle. The model formalizes this force through destination-specific export cutoffs, and the empirical design tests it using pre-release unit values and cross-destination AI intensity, linking global AI diffusion to the composition of U.S. exports.

Third, the paper contributes to research on artificial intelligence and trade by identifying and quantifying a selection channel. Goldfarb and Trefler (2019) argue that AI matters for trade through scale, competition, knowledge creation, and international diffusion, and a growing body of evidence shows that GenAI raises productivity across knowledge work, with the largest gains for workers, tasks, and sellers with lower baseline performance.[5] This pattern motivates the quality-biased channel in our setting. A recent literature studies AI and

[3]On heterogeneous-firm trade and reallocation, see Melitz (2003); Helpman, Melitz, and Yeaple (2004); Bernard, Redding, and Schott (2011); on the uneven incidence of trade shocks, see Bernard, Jensen, and Schott (2006); Autor, Dorn, and Hanson (2013); Caliendo, Dvorkin, and Parro (2019).

[4]On unit values and export prices as quality signals, see Hummels and Klenow (2005); Schott (2004); Hallak and Schott (2011); Manova and Zhang (2012); Kugler and Verhoogen (2012); Bastos, Silva, and Verhoogen (2018); on quality responses to export-market access and competition, see Verhoogen (2008); Brambilla, Lederman, and Porto (2012); Amiti and Khandelwal (2013); Atkin, Khandelwal, and Osman (2017); on quality ladders, quality sorting, and the joint role of quality, markups, and trade costs in export prices, see Khandelwal (2010); Baldwin and Harrigan (2011); Crozet, Head, and Mayer (2012); De Loecker, Goldberg, Khandelwal, and Pavcnik (2016); Fan, Li, Xu, and Yeaple (2020).

[5]On the productivity effects of GenAI and the larger gains at lower baseline performance, see Noy and Zhang (2023); Peng, Kalliamvakou, Cihon, and Demirer (2023); Brynjolfsson, Li, and Raymond (2025); Dell'Acqua, McFowland III, Mollick, Lifshitz, Kellogg, Rajendran, Krayer, Candelon, and Lakhani (2026); Fang, Yuan, Zhang, Donati, and Sarvary (2025). Earlier work shows that machine translation raised platform exports by lowering language frictions (Brynjolfsson, Hui, and Liu, 2019).

trade at the exposure, sector, and firm levels.[6] Closest to our paper, Antoniades, He, Liang, and Xu (2026) show that firm-level AI hiring improves market intelligence and raises export participation. Our mechanism is different. We study a broad GenAI shock and focus on destination-side adoption as a source of low-end foreign competition, so the evidence points to contraction concentrated among initially low-price products in AI-intensive destinations rather than to export expansion through better demand learning. We then quantify this mechanism. The estimated model separates U.S. productivity gains from destination-side selection and lets AI exposure vary along the quality distribution. The decomposition shows that endogenous export cutoffs and quality-biased AI exposure account for the large price response among initially low-price products in high-AI destinations and for the decline in cross-destination unit-value dispersion. The welfare analysis shows that GenAI can raise welfare even when its visible export response is not broad quantity growth, because the technology changes both productivity and export selection.

The rest of the paper proceeds as follows. Section 2 presents descriptive facts on export prices, quantities, and the diffusion of GenAI. Section 3 develops the quality-biased selection model and derives its empirical predictions. Section 4 describes the customs, product-exposure, destination-AI, and industry data. Section 5 presents the reduced-form evidence on export unit values, quantities, destination heterogeneity, and domestic industry outcomes. Section 6 reports robustness checks. Section 7 estimates the model, validates untargeted margins, and uses the estimated structure for mechanism decomposition and welfare analysis. Section 8 concludes.

[6]See Jakubik, Rotunno, and Saini (2025) on AI exposure and bilateral export growth, Sun and Trefler (2023) on AI and digital-services trade, Gozen, Gunes, and Karaca (2025) on AI and firm sales across domestic and foreign markets, and Bonfiglioli, Crinò, Filomena, and Gancia (2025) on comparative advantage in AI-intensive industries.

2 Motivating Facts

This section documents two descriptive patterns that motivate the empirical design and the quantitative model. The first shows that U.S. export prices and quantities moved differently after the release of ChatGPT, a pattern that is difficult to interpret through the standard view of AI as a technology that lowers trade frictions or expands exporter productivity. The second shows that the AI shock diffused rapidly across countries, making destination-market adoption central to the analysis.

1. Export unit values rise and export quantities fall after the ChatGPT release. Figure 1 plots the average log export unit value (Panel A) and the average log export quantity (Panel B) of U.S. products in each month from 2020 through 2025, both normalized to zero in November 2022, the month before the public release of ChatGPT.[7] A product's export unit value is its total export value divided by its total export quantity.

Before the release, the two panels do not move systematically apart: export unit values and quantities show no sustained opposing trend in the months leading up to November 2022. After the release the two series separate and keep diverging through the end of the sample. The smoothed unit-value series rises steadily, reaching about 12 log points above its November 2022 level by December 2025, while the smoothed quantity series falls over the same window to about 24 log points below it. The divergence builds up gradually rather than appearing as a one-month jump, consistent with an adjustment that accumulates as the technology diffuses.

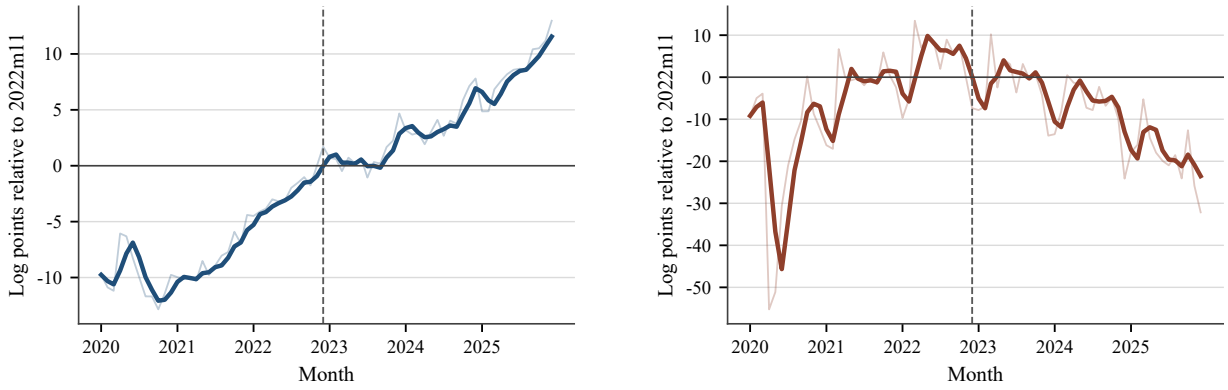
This joint pattern is difficult to reconcile with the way the literature usually models an AI shock. If generative AI mainly raised exporter productivity, it would lower marginal costs and export prices while raising quantities. If it mainly reduced the search, communication, and matching costs of exporting, as in work on digital technologies and trade, it would again

[7]The figure uses a balanced panel of HS10 products with positive export value and quantity in every month of 2020–2025. For each product-month, export value and quantity are aggregated across destinations and the unit value is total value divided by total quantity; each panel plots the equal-weighted mean across products of the corresponding log series. Thin lines show raw monthly means and thick lines trailing three-month moving averages.

make exporting easier and tend to expand quantities. Both views imply that quantities rise, and a uniform productivity gain implies that prices fall; neither generates rising unit values together with falling quantities. A pattern in which the average export price rises while the quantity shipped falls instead points to a change in the composition of what is exported: if generative AI drives lower-priced varieties out of trade, the average unit value of the surviving bundle rises even as total quantity declines. This compositional, selection-based reading is the mechanism that the rest of the paper develops and tests.

Figure 1: Raw Export Outcomes Around the ChatGPT Release

(a): Export unit value **(b): Export quantity**



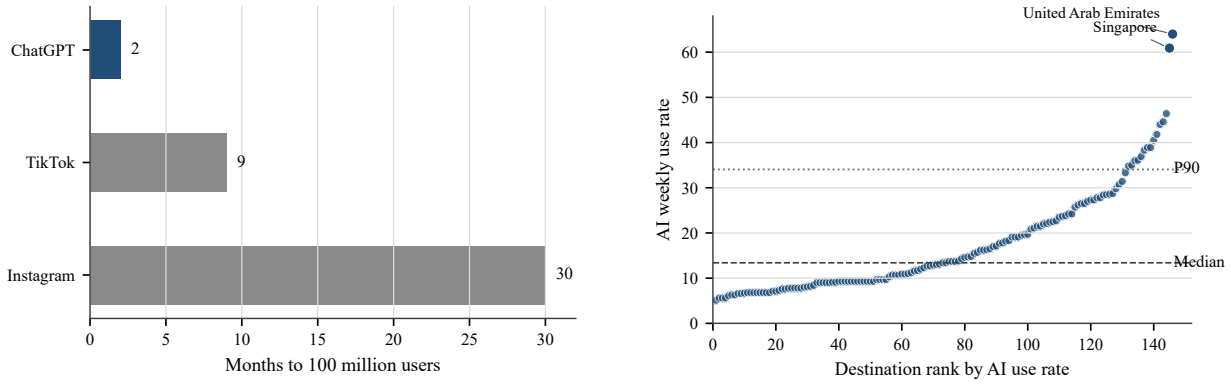
Notes. The figure uses the balanced HS10 product-month panel from 2020m1 through 2025m12, retaining products with positive export value and quantity in every month. Product-month unit values are computed after aggregating export value and quantity across destinations. Panel A plots the equal-weighted mean of log unit values across products; Panel B plots the equal-weighted mean of log quantities. Both series subtract their 2022m11 values. Thin lines are raw monthly values; thick lines are trailing three-month moving averages. The dashed vertical line marks 2022m12.

2. Generative AI diffusion is fast and uneven across destination markets. Figure 2 shows that the relevant technology shock was both rapid and geographically uneven. Panel A places the ChatGPT release in the diffusion path of other consumer technologies. ChatGPT reached 100 million users in about two months, much faster than TikTok or Instagram. This speed matters because it makes the release a plausibly sharp change in the technology environment faced by firms and consumers across countries. Panel B shows that diffusion

was not uniform across destination markets. In the Microsoft AI Diffusion data,[8] the 90th-percentile destination has a weekly use rate more than twice the median, and the highest-use destination is almost five times the median. Thus, after the release, U.S. exporters were not facing a common foreign environment. The same product could be shipped to destinations where GenAI adoption remained limited and to destinations where local firms and consumers were already using the technology intensively. This cross-country dispersion is central for the selection mechanism: destination-side adoption can raise competitive pressure abroad even when product exposure is held fixed.

Figure 2: The Speed and Unevenness of Generative AI Diffusion

(a): Speed of user adoption **(b): Uneven destination use**



Notes. Panel A reports months to 100 million users for ChatGPT, TikTok, and Instagram, based on UBS estimates reported by Hu (2023). Panel B plots destination-level AI diffusion rates and sorts destinations by the country-level use rate. The dashed and dotted horizontal lines in Panel B mark the cross-country median and 90th percentile.

3 Model

We build a quality-augmented heterogeneous-firm model of U.S. exporters serving foreign destinations. Establishments decide whether to serve each destination after drawing a firm

[8]Source: Microsoft AI Economy Institute, *Global AI Adoption in 2025: A Widening Digital Divide*, January 8, 2026, AI Diffusion by Economy data for the second half of 2025, `https://www.microsoft.com/en-us/corporate-responsibility/topics/ai-economy-institute/reports/global-ai-adoption-2025/`.

type. The firm type governs product quality and marginal cost, so selection into export markets changes both the set of active varieties and the unit values observed in customs data. Generative AI enters through two channels. U.S. adoption raises establishment productivity in more exposed products, while destination adoption strengthens low-end foreign competition and reduces residual foreign demand for lower-quality U.S. varieties. These two forces generate predictions for export unit values, export quantities, heterogeneity across pre-period unit-value groups and destination AI intensity, and domestic establishment counts.

The two motivating facts discipline the model. The rise in export unit values without parallel quantity growth pushes against an interpretation based only on a uniform expansion of U.S. productive capacity. The fast but uneven diffusion of GenAI across destination markets makes destination-side adoption a natural margin of adjustment: the same product-level exposure can face different post-release competitive pressure across foreign markets. The model therefore separates U.S. productivity effects from destination-side selection effects and asks which channel can account for the observed price response.

The model builds on the Melitz selection framework (Melitz, 2003) and its quality-augmented extension in Baldwin and Harrigan (2011), with the structural treatment of unit values closest to Johnson (2012). The same product exposure index interacts with two country-level adoption margins: U.S. adoption, common across destinations after the ChatGPT release, and destination AI intensity, measured across foreign markets.

3.1 Environment

Setting and firm heterogeneity. The home country is the United States. Foreign destinations are indexed by $d \in \mathcal{D}$, and the home market is labelled $d = H$. Product categories are indexed by $j \in \mathcal{J}$. Within each product, a continuum of U.S. establishments draws a firm type a from a Pareto distribution with shape k_j and lower bound $\underline{a}_j$,

$$F_j(a) = 1 - \left(\frac{\underline{a}_j}{a}\right)^{k_j}, \qquad a \geq \underline{a}_j, \qquad k_j > \vartheta, \tag{1}$$

with $\vartheta \equiv \eta(\sigma - 1)$ defined below. Quality is a Cobb–Douglas function of the firm type,

$$\xi(a) = a^{1+\eta}, \qquad \eta > 0, \tag{2}$$

following Baldwin and Harrigan (2011). Higher-a establishments sell higher-quality, higher-priced varieties, and the elasticity η governs how strongly observed unit values load on the firm draw. The firm draw a is a quality-cost index rather than a pure productivity index: higher-a varieties carry higher quality but also require more labor per unit of factory-gate output and therefore face higher marginal cost, as shown by equations (10) and (11) below. The maintained restriction $k_j > \vartheta$ ensures finite revenue integrals (Chaney, 2008).

Generative AI exposure. Generative AI exposure is modeled as a multiplicative index that separates three empirically distinct sources of variation: country-time adoption M_{ct}, which captures the diffusion of AI across markets and dates; product exposure s_j, which captures differences in the relevance of language-model tasks across product categories; and a quality-dependent exposure term $\zeta_\eta(\xi(a))$, which allows a given country-product shock to have different intensity across varieties at different points of the quality ladder. For country c, product j, and the quality $\xi(a)$ associated with firm type a, we define

$$\psi_{cjt}(a) \;\equiv\; M_{ct}\, s_j\, \zeta_\eta(\xi(a)), \qquad \zeta_\eta(\xi(a)) \equiv \xi(a)^{-\rho/(1+\eta)} = a^{-\rho}, \qquad \eta > 0, \quad \rho > 0, \tag{3}$$

where $M_{ct} \geq 0$ denotes country-time AI adoption and $s_j \geq 0$ denotes product exposure to language-model tasks. The term $\zeta_\eta(\xi(a))$ is the within-product quality component of exposure. Since $\xi(a) = a^{1+\eta}$, the implied quality elasticity of AI exposure is $-\rho/(1 + \eta)$: higher $\rho/(1 + \eta)$ places relatively more of a given country-product shock on lower-quality varieties. This quality gradient is consistent with evidence that generative AI raises productivity more for lower-baseline tasks and users (Brynjolfsson, Li, and Raymond, 2025; Noy and Zhang, 2023). The same exposure enters two model margins: destination adoption shifts foreign residual demand for U.S. varieties through $\psi_{djt}(a)$, while U.S. adoption shifts U.S. establishment productivity through $\psi_{Ujt}(a)$.

Household preferences and the homogeneous-good numeraire. The U.S. representative household allocates expenditure across a freely traded homogeneous good Z and the $|\mathcal{J}|$ differentiated product categories under Cobb–Douglas preferences,

$$\mathcal{U}_t = \left(C_t^Z\right)^{\beta^Z} \prod_{j\in\mathcal{J}} \left(C_{jt}\right)^{\beta_j}, \qquad \beta^Z + \sum_{j\in\mathcal{J}} \beta_j = 1, \tag{4}$$

with each C_{jt} a CES aggregator of available varieties at elasticity of substitution $\sigma > 1$. Top-tier optimization delivers constant expenditure shares,

$$E_{jt} \equiv \beta_j Y_t, \qquad E_t^Z \equiv \beta^Z Y_t, \tag{5}$$

where Y_t is aggregate U.S. labor income. The homogeneous good Z is produced with the unit-linear technology $Q_{ct}^Z = L_{ct}^Z$ in every country c and is traded at zero iceberg cost; zero profit in the Z sector then pins the producer price of Z to the local wage. We choose Z as the numeraire, $p_t^Z \equiv 1$, which fixes wages in every country at unity and gives $Y_t = L_t$ for the U.S. Foreign wages do not appear elsewhere; the foreign side of the economy enters the model only through the residual demand shifter $\Gamma_{dj,0}$ introduced below.

Variety demand. A U.S. variety of type a sold at price p to destination d at time t faces demand

$$q_{djt}(a,p) = D_{djt}(a)\, \xi(a)^{\sigma-1}\, p^{-\sigma}, \tag{6}$$

with the destination demand shifter $D_{djt}(a)$ specified separately for the home and foreign markets. Home demand follows the CES allocation derived from the inner aggregator of (4),

$$D_{Hjt} = \beta_j L_t\, P_{jt}^{\sigma-1}, \tag{7}$$

where β_j is the Cobb–Douglas expenditure share from (5) and P_{jt} is the CES price index of product j; the home shifter is constant across types within product j. Foreign residual demand for U.S. varieties takes the reduced form

$$D_{djt}(a) = \Gamma_{dj,0}\, \exp\!\big\{-\alpha\, \psi_{djt}(a)\big\}, \qquad d \in \mathcal{D}, \tag{8}$$

in which $\Gamma_{dj,0}$ is a pre-period level that absorbs destination size, the foreign price index for product j, and the strength of competing varieties; $\alpha > 0$ governs how strongly destination AI contracts U.S. residual demand; and $\psi_{djt}(a) = M_{dt}\, s_j\, \zeta_\eta(\xi(a))$ is the destination AI exposure from (3). The exponential form keeps the demand contraction separable from the production block introduced next, and the quality-dependent term inside $\psi_{djt}(a)$ concentrates the contraction on lower-quality U.S. varieties: AI adoption abroad improves lower-quality foreign producers more and therefore tightens competition most strongly against the lower-end U.S. varieties.

Production, marginal cost, and pricing. U.S. generative AI raises establishment productivity through the same primitive ψ. An establishment of type a in product j has effective productivity

$$\Psi_{jt}(a) = \exp\big\{\nu\, \psi_{Ujt}(a)\big\}, \qquad \nu > 0, \tag{9}$$

where $\psi_{Ujt}(a) = M_{Ut}\, s_j\, \zeta_\eta(\xi(a))$ is the U.S. AI exposure from (3) evaluated at $c = U$, $M_{Ut} \geq 0$ is U.S. AI adoption (identified empirically through the post-release indicator Post_t), and $\nu > 0$ governs how strongly U.S. AI raises U.S. establishment productivity. The labor production function is linear,

$$y = A_j\, \frac{\Psi_{jt}(a)}{a}\, \ell, \tag{10}$$

with A_j a product-specific Hicks-neutral component. Under the wage normalization $w_t = 1$ from above, the delivered marginal cost of a type-a firm in destination d is $mc_{djt}(a) = \tau_{dj}\, a/[A_j\, \Psi_{jt}(a)]$, with $\tau_{dj} \geq 1$ an iceberg trade cost. Monopolistic competition with CES demand yields the standard constant-markup price,

$$p_{djt}(a) = \mu\, \tau_{dj}\, \frac{a}{A_j\, \Psi_{jt}(a)}, \qquad \mu = \frac{\sigma}{\sigma - 1}, \tag{11}$$

and substitution into (6) delivers per-firm revenue

$$r_{djt}(a) = D_{djt}(a)\, (\mu \tau_{dj})^{1-\sigma}\, A_j^{\sigma-1}\, \Psi_{jt}(a)^{\sigma-1}\, a^{\vartheta}, \qquad \vartheta = \eta(\sigma - 1). \tag{12}$$

The two AI channels share the common primitive $\psi_{cjt}(a) = M_{ct} s_j \zeta_\eta(\xi(a))$ from (3) but enter revenue (12) with opposite signs: destination AI contracts $D_{djt}(a)$ through the elasticity α, while U.S. AI raises $\Psi_{jt}(a)$ through the elasticity ν. Both effects are stronger for lower-quality varieties when $\rho > 0$.

3.2 Free Entry and Equilibrium

Export selection. Serving destination d in product j requires a per-period fixed labor cost f_{dj}. Operating profit is $\pi_{djt}(a) = r_{djt}(a)/\sigma - f_{dj}$ under the constant-markup rule, and the zero-profit cutoff a^*_{djt} is the smallest a at which profit is non-negative,

$$\frac{1}{\sigma} D_{djt}(a^*_{djt}) \, (\mu \tau_{dj})^{1-\sigma} \, A_j^{\sigma-1} \, \Psi_{jt}(a^*_{djt})^{\sigma-1} \, (a^*_{djt})^{\vartheta} = f_{dj}. \tag{13}$$

Establishments with $a < a^*_{djt}$ do not export to d. For each foreign destination, the cutoff condition is a one-dimensional equation in a^*_{djt},

$$\frac{\Gamma_{dj,0} \, (\mu \tau_{dj})^{1-\sigma} \, A_j^{\sigma-1}}{\sigma} \, (a^*_{djt})^{\vartheta} \, \exp\!\big\{ s_j \, (a^*_{djt})^{-\rho} \big[(\sigma-1)\nu M_{Ut} - \alpha M_{dt} \big] \big\} = f_{dj}. \tag{14}$$

The home cutoff a^{H*}_{jt} is the analogue evaluated at $D_{Hjt} = \beta_j L_t P_{jt}^{\sigma-1}$ and is solved jointly with P_{jt} through the free-entry condition.

Free entry and the establishment count. Potential entrants pay a sunk cost f^E_j in labor units, draw a from F_j, and then decide which destinations to serve. With stationary exit rate δ, free entry equates the expected present value of operating profit across markets to the sunk cost,

$$\delta \, f^E_j = \sum_{d \in \{H\} \cup \mathcal{D}} f_{dj} \int_{a^*_{djt}}^{\infty} \left[\frac{r_{djt}(z)}{r_{djt}(a^*_{djt})} - 1 \right] dF_j(z), \tag{15}$$

where the survivor-conditional revenue ratio inside the integral admits a closed form under the Pareto assumption (1). The home price index then aggregates surviving home-market varieties,

$$P_{jt}^{1-\sigma} = N_{jt} \, \mu^{1-\sigma} \, A_j^{\sigma-1} \, k_j \underline{a}_j^{k_j} \int_{a^{H*}_{jt}}^{\infty} a^{\vartheta - k_j - 1} \, \exp\!\big\{ (\sigma-1)\nu M_{Ut} s_j a^{-\rho} \big\} \, da, \tag{16}$$

and pins the pre-selection mass of active establishments N_{jt}.

Equilibrium. An equilibrium is a collection of prices, quantities, cutoffs, price indices, establishment masses, and labor allocations $\{p_{djt}(a),\, q_{djt}(a),\, a^*_{djt},\, a^{H*}_{jt},\, P_{jt},\, N_{jt}\}_{d\in\mathcal{D},\, j\in\mathcal{J}}$ such that, given $\{M_{Ut}, M_{dt}\}_{d\in\mathcal{D}}$ and product primitives:

1. Firms maximize profits. Active establishments charge the CES markup price in (11); foreign participation cutoffs satisfy (14); and the home participation cutoff satisfies the home-market counterpart of (13);

2. Household demand is optimal. Domestic expenditure follows (5), home demand is summarized by $D_{Hjt} = \beta_j L_t P_{jt}^{\sigma-1}$, and foreign demand for U.S. varieties follows (8);

3. Free entry holds. The expected value of entry equals the sunk entry cost, as in (15);

4. Price indices are consistent with the set of active varieties. The home price index and the pre-selection mass of establishments satisfy (16);

5. Markets clear. Variety-level output equals home and foreign demand, and U.S. labor used in differentiated production, fixed market-access costs, entry costs, and homogeneous-good production sums to the aggregate labor endowment.

Appendix A derives the revenue, unit-value, free-entry, and price-index expressions used in this section. Appendix C gives the equilibrium-solving algebra for arbitrary shock paths.

3.3 Effects of the AI Shock

Local effects around the pre-period steady state. We evaluate the model at the pre-ChatGPT allocation, where $M_{Ut} = M_{dt} = 0$ for every (U, d), $\Psi_{jt}(a) = 1$, and $D_{djt}(\xi) = \Gamma_{dj,0}$ for $d \in \mathcal{D}$. At this point, a change in AI adoption affects a product-destination market through two objects: the product exposure s_j and the type of the marginal exporter, summarized by $(a^*_{dj,0})^{-\rho}$. The response is therefore larger for products with greater language-model exposure and for markets in which the marginal exporter lies lower on the quality ladder.

Proposition 1 (AI adoption, selection, and export unit values)**.** *At first order around the pre-period steady state, the foreign export cutoff response is*

$$d\log a^*_{djt} = \frac{s_j\,(a^*_{dj,0})^{-\rho}}{\vartheta}\big[\,\alpha\,dM_{dt} - (\sigma-1)\nu\,dM_{Ut}\,\big], \tag{17}$$

and the export unit-value and per-firm export quantity responses are

$$d\log p^x_{djt} = s_j\,(a^*_{dj,0})^{-\rho}\big[\,K^d_j\,dM_{dt} - K^U_j\,dM_{Ut}\,\big], \tag{18}$$

$$d\log Q^x_{djt}\,\big|_{dN=0} = s_j\,(a^*_{dj,0})^{-\rho}\big[\,-K^{Q,d}_j\,dM_{dt} + K^{Q,U}_j\,dM_{Ut}\,\big], \tag{19}$$

with $K^d_j > 0$, $K^U_j > 0$, $K^{Q,d}_j > 0$, *and* $K^{Q,U}_j > 0$. *Appendix B derives these coefficients.*

The signs follow from the structure of (14). Destination AI raises the export cutoff because the residual demand shifter $D_{djt}(\xi(a^*))$ contracts and the marginal exporter loses the variable profit needed to cover the fixed cost. The cutoff move removes lower-quality U.S. varieties from the destination, and the conditional mean unit value of surviving exporters rises. U.S. AI lowers the cutoff because $\Psi_{jt}(a^*) > 1$ raises variable profit at every type, admitting lower-quality entrants and pulling the conditional mean unit value down. The two channels therefore move export unit values in opposite directions. The model predicts a positive aggregate unit-value response only when the destination-demand selection channel is stronger than the U.S. supply-side entry channel in the relevant export portfolio.

The multiplier $(a^*_{dj,0})^{-\rho}$ determines how the effect varies with the pre-period export cutoff in a product-destination market. At the baseline, the model implies

$$p^x_{dj,0} = \frac{\mu\tau_{dj}}{A_j}\,\frac{k_j - \vartheta + 1}{k_j - \vartheta}\,a^*_{dj,0}.$$

Holding the cost shifter fixed, lower baseline unit values correspond to lower cutoffs and therefore to larger values of $(a^*_{dj,0})^{-\rho}$. The model then predicts larger responses, in absolute value, for lower-cutoff markets. Destination AI raises export unit values and lowers quantities more strongly in those markets, while U.S. AI lowers export unit values and raises quantities more strongly. Higher destination adoption M_{dt} strengthens the destination-demand channel

through the product $s_j(a^*_{dj,0})^{-\rho}M_{dt}$.

Proposition 2 (U.S. AI and entry)**.** *The first-order response of the pre-selection mass of U.S. establishments satisfies*

$$\begin{aligned} d\log N_{jt} &= s_j\left[B^U_j\, dM_{Ut} - \sum_{d\in\mathcal{D}} B^D_{dj}\, dM_{dt}\right], \\ B^U_j &= (\sigma-1)\nu\left[\sum_{d\in\mathcal{D}} \omega^N_{dj,0}(a^*_{dj,0})^{-\rho}\Lambda^N_j - \Theta^H_j(a^*_{Hj,0})^{-\rho}\right], \\ B^D_{dj} &= \alpha\,\omega^N_{dj,0}\Lambda^N_j(a^*_{dj,0})^{-\rho} > 0, \\ \Lambda^N_j &= \frac{k_j(k_j-\vartheta)}{\vartheta(k_j+\rho-\vartheta)}\frac{R^X_{j,0}}{R^H_{j,0}} > 0, \qquad \Theta^H_j = \frac{\rho(k_j-\vartheta)}{\vartheta(k_j+\rho-\vartheta)} > 0. \end{aligned} \tag{20}$$

where $R^H_{j,0}$ is baseline home-market revenue, $R^X_{j,0} \equiv \sum_{d\in\mathcal{D}} R^x_{dj,0}$ is baseline foreign export revenue, and $\omega^N_{dj,0} = R^x_{dj,0}/R^X_{j,0}$. The terms Λ^N_j and Θ^H_j are equilibrium pass-through coefficients from the free-entry system. The former maps foreign-market profit changes into establishment mass, while the latter maps the home-market price-index adjustment into establishment mass. Positivity follows from $k_j > \vartheta$, $\rho > 0$, and positive baseline home and foreign revenues.

Equation (20) decomposes the entry incentive into domestic and foreign AI components. U.S. adoption raises U.S. productivity and therefore raises the value of operating in export markets. This export-market force is the first term inside B^U_j, summed over destinations and weighted by their baseline revenue shares. U.S. adoption also affects the domestic equilibrium. The home-market term $\Theta^H_j(a^*_{Hj,0})^{-\rho}$ is the offset generated by the home price index and the free-entry condition. Hence B^U_j is a net domestic-adoption coefficient, not a primitive parameter. It is positive only when the export-market profit force dominates the home-market equilibrium offset.

Foreign adoption works through a different margin. Higher M_{dt} improves foreign producers in destination d, especially lower-quality producers. This weakens the residual demand faced by U.S. exporters in that market. The strength of this effect is proportional to destination d's baseline importance, $\omega^N_{dj,0}$, and to the baseline cutoff weight $(a^*_{dj,0})^{-\rho}$. These

pre-period objects indicate where the establishment margin is most exposed to foreign AI adoption; they are not themselves new shocks.

The model therefore delivers a clear prediction for domestic producer mass. Domestic AI can raise the number of active U.S. establishments when the net coefficient B_j^U is positive. Destination-market AI lowers the number of active U.S. establishments through the negative foreign-adoption terms. With both shocks present, the establishment response is governed by the balance between the domestic productivity force and the foreign competitive force.

The propositions provide the mapping from the model to the empirical design. The empirical counterparts of the model objects are defined in the next section. In that mapping, the product exposure term s_j corresponds to product-level GenAI exposure; U.S. adoption is captured by the post-ChatGPT period; and cross-destination variation in destination-side adoption is proxied by pre-release destination AI intensity. The cutoff term $(a^*_{dj,0})^{-\rho}$ is not observed directly, but the model links it to pre-release export unit values, with lower pre-release unit values corresponding to lower export cutoffs and therefore greater exposure to export selection. The empirical analysis therefore focuses on three predictions from the model, namely product-level price and quantity responses, heterogeneity by pre-release unit values and destination AI intensity, and domestic establishment responses.

4 Data Description

4.1 Data

We construct four samples spanning product-, product-destination-, industry-, and industry-destination-level observations from monthly U.S. customs records and annual federal statistical data. Each sample is matched to a measure of industry-level GenAI exposure.

HS10 monthly export records. Monthly product-destination shipments come from the U.S. Census Bureau's Schedule B records, which report export value, quantity, and destination country for each HS10 product-destination-month observation. We restrict the sample to

records with identified destination countries and require positive export value and positive quantity for any price calculation. The monthly data span January 2020 through December 2025, covering 35 months before and 37 months after the public release of ChatGPT on November 30, 2022.

Industry-destination related-party trade. Annual domestic exports by four-digit NAICS industry and destination, classified by related-party status, come from the Census Related-Party Trade Data. The series is published at the six-digit NAICS level, and we aggregate dollar values across constituent six-digit codes within each four-digit parent industry to align with the four-digit dimensionality of the GenAI exposure measure. We restrict the sample to identified destinations and valid four-digit NAICS industries, and separate each industry-destination-year observation into related-party, arm's-length, and unreported-status components. These data are available from 2020 through 2024.

Industry employment and establishment dynamics. Industry-level employment, establishment counts, and average establishment size come from the Quarterly Census of Employment and Wages (QCEW), aggregated from QCEW's native six-digit NAICS publication to four-digit NAICS by summing employment and establishment counts across constituent six-digit codes. Establishment entry counts come from the Business Dynamics Statistics (BDS), which publishes entry directly at the four-digit NAICS level. Both series span 2020 through 2025 and are observed at the four-digit NAICS industry by calendar year level.

Treatment and destination measures. Industry-level GenAI exposure is constructed from the AI Industry Exposure dataset of Felten, Raj, and Seamans (2021) and is described in Section 4.3. Destination-country AI intensity is measured by the OECD/LinkedIn AI hiring rate and AI talent concentration, and is described in Section 4.5.

4.2 The Generative AI Shock

Our empirical strategy treats the public release of ChatGPT on November 30, 2022 as a discrete, economy-wide shock. The shock applies to production tasks that can be performed

by large language models. Transformer architecture and related large language model research had developed rapidly since 2017. Before ChatGPT, however, firms and households did not have access to a comparable general-purpose generative AI application. Its release made a frontier language model directly accessible through a web browser. It converted an advanced research capability into a widely usable consumer and business technology. Because the product became available to firms across industries at the same point in time, the release provides a natural dated shock. This timing motivates our use of December 2022, $t = 2022\text{m}12$, as the onset of treatment rather than modeling generative AI exposure as a gradually evolving shock.

The timing of the release matters for identification. ChatGPT reached one million users within five days and one hundred million monthly active users within two months.[9] U.S. adoption was faster than early internet adoption, and cross-country evidence documents GenAI use across more than 100 economies soon after the release (Bick, Blandin, and Deming, 2024; Ali, Mustafa, and Aysan, 2025). These facts make December 2022 a natural common break in both U.S. exporters' own technology environment and the technology environment of their destination markets. Financial-market evidence supports the same timing assumption. Eisfeldt, Schubert, and Zhang (2023) document that references to generative AI in corporate earnings calls were essentially absent before the fourth quarter of 2022 and increased sharply thereafter.[10]

The selection interpretation comes from the incidence of the technology across producers. Experimental and workplace studies find larger gains for lower-performing or less experienced workers in writing and customer service tasks (Noy and Zhang, 2023; Brynjolfsson, Li, and Raymond, 2025). Field evidence from online retail finds larger gains among smaller sellers (Fang, Yuan, Zhang, Donati, and Sarvary, 2025). In our setting, this evidence motivates

[9]Eisfeldt, Schubert, and Zhang (2023) note that this diffusion rate was substantially faster than that of earlier consumer technologies such as the World Wide Web, WhatsApp, Instagram, and Twitter, each of which took between two-and-a-half and seven years to reach comparable scale.

[10]They also show that a long-short portfolio sorted on industry-level generative AI exposure earns an abnormal return of nearly five percent in the ten trading days following November 30, 2022, with no comparable pre-trend in the preceding window and no reversal over the subsequent two months.

the focus on destination-market adoption. If GenAI has larger effects among lower-baseline-performance foreign sellers, the release changes the competitive environment facing U.S. exporters. It can also shift the set of U.S. varieties observed in trade.

4.3 Identifying Industries and Products Exposed to Generative AI

Industry-level GenAI exposure comes from the AI Industry Exposure (AIIE) dataset of Felten, Raj, and Seamans (2021). They construct the underlying occupation-level AI Occupational Exposure (AIOE) in three steps. First, they identify ten well-developed AI applications recognized by the Electronic Frontier Foundation as having experienced meaningful scientific progress, including language modeling, image recognition, image generation, translation, and speech recognition. They then link each application to the 52 occupational abilities in O*NET using application-ability relatedness scores collected from Amazon Mechanical Turk surveys, and sum these scores across applications to obtain an ability-level AI exposure. Second, for each six-digit SOC occupation they weight the ability-level exposure by the prevalence and importance of each ability within that occupation as recorded in O*NET, then scale by the occupation's total ability breadth so that the resulting AIOE measures *relative* rather than *aggregate* exposure. Third, the occupation-level AIOE is aggregated to four-digit NAICS industries (AIIE) using 2019 employment shares from the Bureau of Labor Statistics. We use the language modeling application of this methodology, which Felten, Raj, and Seamans (2021) subsequently released as part of their generative AI dataset extension and which captures industry-level exposure to language-model-based generative AI in particular.[11]

Product-level GenAI exposure extends the four-digit NAICS industry measure to HS10 products through the U.S. Census Bureau's six-digit NAICS to HS10 Schedule B concordance. The full assignment chain therefore runs from four-digit NAICS AIIE to six-digit NAICS through the four-digit parent and from six-digit NAICS to HS10 through the Sched-

[11]Data and replication code are publicly available at `https://github.com/AIOE-Data/AIOE`, and the language modeling scores are in the file `Language Modeling AIOE and AIIE.xlsx`.

ule B concordance. The resulting product-level measure, denoted GenAI_p, is observed in a product-month sample of 5,788 HS10 products and 416,736 observations and in a product-destination-month sample of 5,779 products and 7,398,106 trade observations.

4.4 Outcome Variables

The empirical analysis uses four units of observation: HS10 product-month, HS10 product-destination-month, four-digit NAICS industry-destination-year, and four-digit NAICS industry-year. Outcome variables are constructed separately at each level to capture price, quantity, trade-composition, and industrial-organization responses to generative AI exposure. The underlying customs data report, for each HS10 product shipped to each destination country in each month, an export value in U.S. dollars and an export quantity. All continuous outcomes are winsorized at the first and ninety-ninth percentiles of their pooled distribution to limit the influence of reporting outliers in customs and federal statistical records.

Product-month outcomes. The price response to generative AI exposure is measured at the HS10 product-month level using two complementary variables. The first is the log aggregate export price, defined as the natural logarithm of the product's total export value across destinations divided by its total export quantity in the same month. The second variable is the cross-destination dispersion of bilateral log unit values, computed as the standard deviation of destination-specific log unit values for the product in the month. It is evaluated only for product-months with at least four active destinations. Taking logs before computing the standard deviation makes the measure scale-invariant in the product's overall price level and therefore comparable across HS10 products with very different baseline unit values, in line with the trade-prices literature that uses within-product price dispersion across markets as evidence of destination-specific quality positioning and demand matching (Schott, 2004; Manova and Zhang, 2012). Two additional product-month variables, the mean log bilateral quantity and the mean log bilateral value, are computed as cross-destination averages of the log of one plus the bilateral quantity and the log of one plus the bilateral

value, a transformation that retains product-months with sparse bilateral shipments.

Product-destination-month outcomes. Outcomes on product-destination level use the same shipment records at the disaggregated HS10 product by destination country by calendar month level. The log export unit value is the natural logarithm of export value divided by export quantity for the product-destination-month observation, and is defined only when both are strictly positive. The log export quantity and the log export value are constructed as the log of one plus the corresponding raw value. These outcomes separate within-product margin adjustments from the cross-destination composition effects that are absorbed when destinations are aggregated to the product-month level.

Industry-year outcomes. The industrial-organization response to generative AI is measured by four annual outcomes at the four-digit NAICS industry level: industry employment, establishment count, establishment entry, and average employment per establishment. Employment, establishment count, and average establishment size come from the Quarterly Census of Employment and Wages (QCEW), aggregated from its native six-digit NAICS publication to four-digit NAICS by summing across constituent six-digit codes. Establishment entry counts come from the Business Dynamics Statistics (BDS) at the four-digit NAICS level, which BDS publishes directly. Average employment per establishment is computed as the ratio of QCEW employment to QCEW establishment count within the four-digit industry.

Industry-destination-year trade composition. Industry-destination-year trade outcomes use the Related-Party Trade Data and are observed in U.S. dollars at the four-digit NAICS industry by destination country by calendar year level. Two dollar-value components are recorded: related-party exports and arm's-length exports. Each is entered as the log of one plus the dollar value to retain observations with zero shipments in one or more components.

Pre-release price groups. For heterogeneity analysis, products are divided into low-price and high-price groups using each HS10 product's average aggregate export unit value over January–November 2022, before the public release of ChatGPT. The cutoff is the median

across HS10 products within their HS4 categories in that pre-release distribution.

4.5 Destination-Country AI Capacity

Cross-country variation in pre-release capacity to absorb generative AI is captured by two indicators of AI labor-market conditions, jointly published by the OECD AI Policy Observatory and LinkedIn through the OECD.AI live data portal.[12] The two indicators capture complementary margins of pre-release AI exposure: destination-country demand for AI workers and the depth of the country's resident pool of AI talent.

The first indicator, the AI hiring rate, measures destination-country demand for AI labor. For each country, the indicator is defined as the share of newly hired LinkedIn members who report AI-related skills, smoothed by a trailing twelve-month moving average.[13] We extract the value for November 2022, the last full month before the public release of ChatGPT, so that the indicator is not mechanically affected by the post-release surge in AI labor demand.[14]

The second indicator, AI talent concentration, measures destination-country supply of AI labor. It is defined as the share of LinkedIn members in a country who are classified as AI talent.[15] We use the 2021 cross-section, the most recent annual snapshot before the ChatGPT release, to measure pre-release AI talent supply.[16]

For each indicator, the cross-country median splits destinations into high- and low-capacity groups, and bilateral price, quantity, and trade-composition responses are estimated

[12]The indicators are produced by the LinkedIn Economic Graph team from anonymized and aggregated profile information for more than one billion LinkedIn members worldwide and are publicly distributed through OECD.AI. Construction details are documented at `https://oecd.ai/en/partner-data-methodological-note` and `https://oecd.ai/en/linkedin`.

[13]Specifically, the numerator is the number of LinkedIn members with AI-related skills on their profile who started a new position with a new employer in the reference month, and the denominator is the total number of LinkedIn members in the country who started a new position with a new employer in the same month. The AI hiring rate series is publicly visualized at `https://chart.oecd.ai/19570-ai-hiring-over-time`.

[14]OECD.AI reports the series only for countries with at least 100,000 LinkedIn members, LinkedIn labor-force coverage of at least 40 percent, and at least ten AI hires in the reference month.

[15]LinkedIn classifies a member as AI talent if the member explicitly lists at least two AI engineering skills on their profile or performs an AI occupation. The cross-country AI talent concentration series is publicly visualized at `https://chart.oecd.ai/12935-ai-talent-concentration-by-country`.

[16]Country coverage spans all G20 and OECD economies and several additional reporting jurisdictions, including Cyprus, Hong Kong (China), Romania, Singapore, the United Arab Emirates, and Uruguay.

separately within each group.

4.6 Specification

We estimate a product-level difference-in-differences specification that compares the evolution of export outcomes before and after the public release of ChatGPT across HS10 products with different exposure to generative AI. Specifically, we estimate

$$Y_{pt} \;=\; \beta \left(\text{GenAI}_p \times \text{Post}_t\right) + \alpha_p + \gamma_t + \varepsilon_{pt}. \tag{21}$$

where p indexes HS10 products and t indexes year-months from 2020m1 through 2025m12. The outcome Y_{pt} is the log aggregate export unit value of product p in year-month t. The post-release indicator is defined as $\text{Post}_t = \mathbf{1}(t \geq 2022\text{m}12)$, where the post period begins with the public release of ChatGPT. The treatment intensity, GenAI_p, measures product-level exposure to generative AI.[17] Product fixed effects, α_p, absorb time-invariant differences across HS10 products, while year-month fixed effects, γ_t, absorb aggregate shocks common to all products. Standard errors are clustered at the HS10 product level.

The coefficient of interest, β, captures whether products with higher GenAI exposure experienced differential changes in export unit values after the public release of ChatGPT, relative to products with lower exposure.

[17]The exposure measure is constructed from inputs that pre-date the public release of ChatGPT. The underlying AI Industry Exposure measure combines O*NET occupational ability profiles, Amazon Mechanical Turk assessments of the relatedness between ten AI applications and fifty-two occupational abilities, and 2019 BLS occupational employment shares at the four-digit NAICS level.

4.7 Descriptive Statistics

Table 1: Descriptive Statistics

Variable	*N*	Mean	SD	Min	Max
Panel A. HS10 product-month panel					
GenAI product exposure	416,736	−0.638	0.655	−1.740	1.398
Post-period indicator	416,736	0.514	0.500	0.000	1.000
High-price product indicator	387,648	0.487	0.500	0.000	1.000
Ln Aggregate Export Unit Value	366,305	3.466	2.586	0.235	11.557
S.D. of Ln Export Unit Value	336,329	0.980	0.650	0.002	2.886
Mean Ln Export Quantity	366,305	6.674	2.725	0.896	12.918
Mean Ln Export Value	366,545	10.435	1.136	8.153	14.482
Panel B. HS10 product-destination-month panel					
High AI hiring destination indicator	3,989,544	0.515	0.500	0.000	1.000
High AI talent destination indicator	4,006,955	0.528	0.499	0.000	1.000
Ln Export Unit Value	7,380,045	4.396	2.613	0.315	10.775
Ln Export Quantity	7,398,106	6.472	3.349	0.693	14.633
Panel C. Domestic industry-year panel					
Industry employment	1,112	392,980	626,000	0	3,585,895
Establishments	1,112	31,851	65,697	0	415,769
Establishment entries	896	2,759	5,288	7	31,122
Average employment per establishment	1,039	28.0	30.9	2.4	173.6
Panel D. Industry-destination-year trade panel					
Ln Related-Party Exports	58,767	7.167	6.819	0.000	20.267
Ln Arm's-Length Exports	58,767	12.721	4.713	0.000	20.720

Notes. The panels follow the order of the empirical analysis in Section 5: HS10 product–month outcomes, HS10 product–destination–month outcomes, domestic four-digit NAICS industry–year outcomes, and four-digit NAICS industry–destination–year trade outcomes. Price, quantity, and value outcomes are in logarithms unless otherwise indicated. Continuous outcomes are winsorized at the 1st and 99th percentiles. In Panel A, Ln Aggregate Export Unit Value is computed from national totals; S.D. of Ln Export Unit Value pools destination-level unit values across destinations within each product and month; Mean Ln Export Quantity and Mean Ln Export Value are cross-destination averages of destination-level $\ln(\text{quantity}+1)$ and $\ln(\text{value}+1)$, as in Section 4.4.

Table 1 reports summary statistics for the four estimation samples. First, product-level exposure to generative AI exhibits substantial cross-sectional variation across HS10 products. The mean value of pre-release GenAI is −0.638, with a standard deviation of 0.655 and a range from −1.740 to 1.398. This dispersion indicates that products differ meaningfully in

their pre-existing exposure to generative AI, providing the variation across products used in the difference-in-differences specifications.

Second, export outcomes vary considerably across products and destinations. At the product-month level, Ln Export Unit Value has a mean of 3.466 and a standard deviation of 2.586, indicating substantial heterogeneity in average export unit values across HS10 products. Within-product cross-destination price dispersion, measured by S.D. of Ln Export Unit Value, has a mean of 0.980 and a standard deviation of 0.650, suggesting that export unit values vary not only across products but also across destinations within the same product. At the product–destination–month level, Ln Export Unit Value has a mean of 4.396 and a standard deviation of 2.613, while Ln Export Quantity has a mean of 6.472 and a standard deviation of 3.349.

Third, the domestic industry–year panel shows wide dispersion in organizational scale. Industry employment has a mean of 392,980 and a standard deviation of 626,000, establishment counts have a mean of 31,851 and a standard deviation of 65,697, and establishment entries have a mean of 2,759 and a standard deviation of 5,288. Average employment per establishment has a mean of 28.0 and a standard deviation of 30.9.

Fourth, the industry–destination–year panel shows large differences between arm's-length and related-party trade. Ln Arm's-Length Exports has a mean of 12.721 and a standard deviation of 4.713, whereas Ln Related-Party Exports has a mean of 7.167 and a standard deviation of 6.819.

5 Empirical Results

We take these predictions to the data as follows. The first set of tests corresponds to Proposition 1. We estimate the average product-level export-price response to GenAI exposure, examine quantities and cross-destination unit-value dispersion, and test the selection mechanism more directly by comparing low- and high-pre-release unit-value products and by allowing the response to vary with destination AI intensity. The establishment and gross-entry

regressions correspond to Proposition 2. They test the domestic side of the same global AI shock: U.S. adoption can raise the value of entry in exposed industries, while foreign adoption can weaken that incentive through stronger selection in export markets. We also report the arm's-length and related-party split as an additional check on the transactional organization of export sales.

5.1 Aggregate Export Unit Value

Table 2 reports difference-in-differences estimates of Equation (21), with the logarithm of the aggregate HS10 export unit value as the dependent variable. All specifications absorb HS10 product and year–month fixed effects. Column (2) additionally allows each HS4 category to follow its own linear time trend. Standard errors are clustered at the HS10 product level.

The estimated difference-in-differences coefficient is positive and statistically significant at the 1% level in both specifications. The baseline estimate in Column (1) is 0.0508. After allowing for HS4-specific linear trends in Column (2), the coefficient declines modestly to 0.0404, suggesting that the result is not driven by differential trends across broader product categories. In terms of economic magnitude, a one-standard-deviation increase in GenAI exposure is associated with a 3.3% larger post-shock increase in aggregate export unit values in the baseline specification and a 2.6% larger increase after controlling for HS4-specific trends.

Taken together, the estimates indicate that products with greater GenAI exposure experienced larger increases in export unit values following the public release of ChatGPT. This establishes the aggregate price response predicted by Proposition 1 when the destination-side selection force dominates.

Table 2: Effects of GenAI Exposure on Aggregate Export Unit Values

	(1)	(2)
	Ln Aggregate Export Unit Value	
	Baseline	With HS4 trend
GenAI × Post	0.0508***	0.0404***
	(0.0077)	(0.0079)
Observations	366,286	366,286
R^2	0.9639	0.9649
HS10 product FE	Yes	Yes
Month FE	Yes	Yes
HS4 × month trend	No	Yes
Clustered SE (HS10)	Yes	Yes

Notes. The table displays difference-in-differences estimates from the HS10 product–month panel (2020m1–2025m12) in which the dependent variable is aggregate export unit value. Aggregate export unit value is the logarithm of total export value divided by total export quantity, aggregated across destinations within each HS10 product and month. The reported coefficient is $\text{GenAI}_p \times \text{Post}_t$. GenAI_p is pre-release GenAI product exposure constructed from occupational task content and fixed pre-period industry employment weights, as in Section 4. Post_t equals one for months on or after the public release of ChatGPT (2022m12) and zero otherwise. Column (1) includes HS10 product and year–month fixed effects; column (2) adds HS4-specific linear month trends. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

5.2 Event-Study Evidence

To assess the parallel-trends assumption underlying the baseline difference-in-differences estimates and to trace the dynamics of the response, we estimate an event-study analog of Equation (21) that replaces the post-release indicator with a full set of leads and lags relative to the ChatGPT release month:

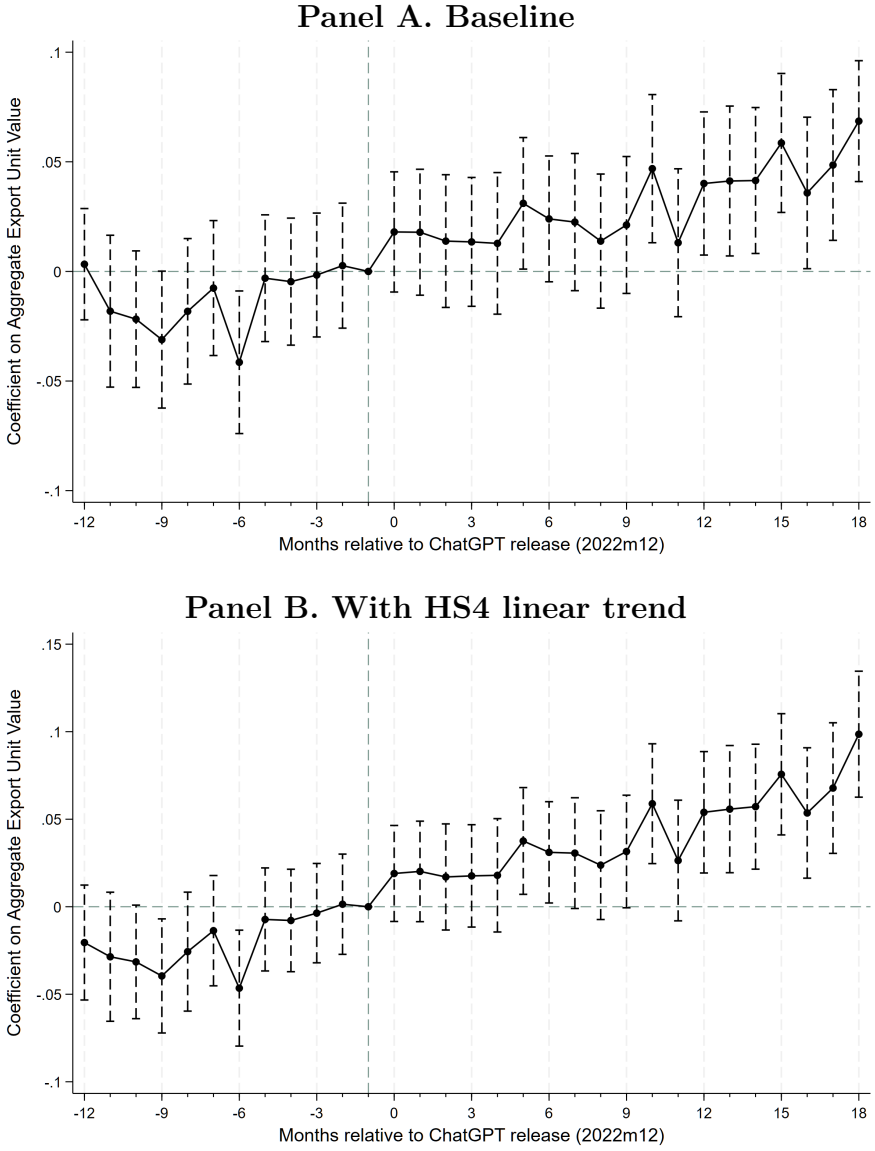
$$Y_{pt} = \sum_{k \neq -1} \beta_k \left(\text{GenAI}_p \times \mathbf{1}\{\tau_t = k\} \right) + \alpha_p + \gamma_t + \varepsilon_{pt}, \tag{22}$$

where $\tau_t = t - 2022\text{m12}$ denotes event time in months relative to the ChatGPT shock. The index k ranges from -12 to 18, and $k = -1$ (November 2022) is omitted as the reference month. The pre-shock coefficients $\{\beta_k\}_{k<-1}$ provide the empirical analog of the parallel-trends assumption, while the post-shock coefficients $\{\beta_k\}_{k\geq 0}$ trace the monthly dynamics of the response. Standard errors are clustered at the HS10 product level.

Figure 3 plots the estimated coefficients β_k together with 95% confidence intervals, in the baseline specification (Panel A) and after additionally absorbing HS4-specific linear month trends (Panel B). The pre-shock coefficients are small and statistically insignificant, with no evidence of differential trends between products with higher and lower GenAI exposure pre-shock. This pattern supports the parallel-trends assumption underlying our identification strategy.

After the ChatGPT shock, the coefficients rise gradually and become significantly positive over the post-shock window. These dynamics indicate that the increase in export unit values is concentrated after the public release of ChatGPT. Overall, the event-study evidence corroborates the baseline difference-in-differences estimates and supports the timing of the model shock.

Figure 3: The timing of aggregate export unit value increases after GenAI exposure



Notes. The figure displays event-study coefficients from the HS10 product–month panel (2020m1–2025m12) for aggregate export unit value. Each coefficient is the interaction between product-level GenAI exposure and an indicator for the calendar month relative to 2022m12; 2022m11 is the omitted reference month. Panel A includes HS10 product and year–month fixed effects; Panel B additionally absorbs HS4-specific linear month trends. Shaded bands show 95% confidence intervals based on standard errors clustered at the HS10 product level.

5.3 Quantity and Dispersion Responses

The aggregate unit-value increase documented in subsection 5.1 can arise from either of the two AI channels in Proposition 1. The proposition gives opposite predictions for export unit

values and quantities. Destination adoption raises export cutoffs, removes lower-quality U.S. varieties from foreign markets, raises observed export unit values, and lowers quantities. U.S. adoption lowers export cutoffs by raising U.S. productivity, which lowers observed export unit values and raises quantities.

To distinguish between these channels, we re-estimate the baseline difference-in-differences specification using mean log export quantity. We also report the cross-destination standard deviation of log export unit values and mean log export value to describe how the aggregate unit-value response is related to destination dispersion and export revenue.

Table 3 reports the estimates. The estimates for mean export quantities are economically small and statistically insignificant in both specifications. This near-zero average response is itself informative for the model. By Proposition 1, U.S. adoption raises export quantities by improving exporter productivity, while destination adoption lowers them by tightening foreign export selection. A shock operating through only one margin would therefore move average quantities, upward under U.S. adoption alone and downward under destination adoption alone. The statistically insignificant average response instead indicates that both forces are present and roughly offset on the quantity margin. Together with the significant unit-value increase in subsection 5.1, the limited quantity response reflects a global AI shock with both a U.S. productivity component and a destination selection component, and the quantitative model uses the aggregate quantity response to discipline their relative strength.

The dispersion and export-value estimates provide additional aggregate evidence for this interpretation. Proposition 1 implies heterogeneous adjustment along the pre-release unit-value distribution. Stronger selection against lower-unit-value varieties reduces their quantities, while the U.S. productivity channel can support expansion among higher-unit-value varieties. Since lower-unit-value varieties tend to start from larger shipment volumes and higher-unit-value varieties from smaller ones, this reallocation can raise the average observed unit value and compress the cross-destination unit-value distribution. The implication for export value is ambiguous ex ante: the contraction of larger, lower-unit-value shipments

works against export value, while expansion among higher-unit-value shipments works in the opposite direction. Columns (5)–(6) report the net of these two forces. Thus, Table 3 is consistent with the selection mechanism in Proposition 1: average quantities change little, unit values rise, dispersion falls, and the value response depends on the relative strength of contraction at the lower end and expansion at the higher end. The next subsections examine these heterogeneous responses directly.

Table 3: Effects of GenAI Exposure on Dispersion, Quantities, and Export Values

	(1)	(2)	(3)	(4)	(5)	(6)
	S.D. of Ln Export Unit Value		Mean Ln Export Quantity		Mean Ln Export Value	
	Baseline	HS4 trend	Baseline	HS4 trend	Baseline	HS4 trend
GenAI × Post	−0.0096***	−0.0063*	−0.0076	0.0002	0.0245***	0.0280***
	(0.0033)	(0.0037)	(0.0088)	(0.0095)	(0.0058)	(0.0066)
Observations	287,348	287,348	366,286	366,286	366,526	366,526
R^2	0.8277	0.8312	0.9247	0.9261	0.7507	0.7545
HS10 product FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × month trend	No	Yes	No	Yes	No	Yes
Clustered SE (HS10)	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The table displays difference-in-differences estimates from the HS10 product–month panel. Columns (1)–(2) use the standard deviation of export unit values across destinations; columns (3)–(4) use mean export quantity; and columns (5)–(6) use mean export value. Aggregate export unit value is constructed from national totals, as in Table 2. The dispersion measure pools destination-level unit values within each product and month and is defined for product–months with more than three active destinations. Mean quantity and mean value are cross-destination averages of $\ln(\text{quantity}+1)$ and $\ln(\text{value}+1)$. Odd-numbered columns are baseline specifications with HS10 product and year–month fixed effects; even-numbered columns add HS4-specific linear month trends. GenAI_p and Post_t are defined in Section 4. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

5.4 Heterogeneous Export Responses by Pre-Release Product Unit Value

Proposition 1 predicts that the U.S. export unit-value and quantity responses to either AI channel are larger in absolute value for product-destination markets with lower pre-period export cutoffs. The cutoff is not directly observed, so we use pre-release unit values as the empirical counterpart of the cutoff-weighting term. In the model, holding cost shifters fixed, lower baseline unit values imply lower cutoffs and larger values of $(a^*_{dj,0})^{-\rho}$. If the foreign-adoption selection channel is stronger in the data, products with lower pre-release average unit values should therefore display larger post-release unit-value increases and, in specifications where quantity effects are visible, larger quantity declines than products with higher pre-release average unit values. We test this prediction by splitting HS10 products at the HS4-specific median of average export unit value within each HS4 category over the pre-shock period January–November 2022 and re-estimating the baseline difference-in-differences specification separately within each subsample.

Table 4 reports the estimates, and Figure 4 plots the same coefficients with 95% confidence intervals, making the asymmetry between low- and high-unit-value subsamples visually transparent. The export unit-value response in columns (1) and (2) is positive in both subsamples but substantially larger and more precisely estimated among low-unit-value products. In the baseline specification (Panel A), the estimated coefficient is 0.0524 in the low-unit-value subsample, significant at the 1% level, and 0.0169 in the high-unit-value subsample, significant at the 5% level. Adding HS4-specific linear month trends in Panel B preserves the asymmetry. The low-unit-value coefficient declines to 0.0298 but remains significant at the 1% level, while the high-unit-value coefficient falls to 0.0143 and retains only marginal significance at the 10% level. A randomization test rejects equality of the low- and high-unit-value coefficients in both panels, with empirical p-values of 0.000 and 0.020.

The export quantity response in columns (3) and (4) displays the same asymmetry with the opposite sign in the baseline specification but is sensitive to HS4-specific linear month

trends. In Panel A, the estimated coefficient is -0.0296 in the low-unit-value subsample, significant at the 5% level, while the high-unit-value coefficient of 0.0117 is small and statistically insignificant, and a randomization test rejects equality of the two with an empirical p-value of 0.000. In Panel B with HS4-specific linear month trends, both coefficients become close to zero, 0.0014 in the low-unit-value subsample and 0.0027 in the high-unit-value subsample, and the randomization test no longer rejects equality, with an empirical p-value of 0.910.

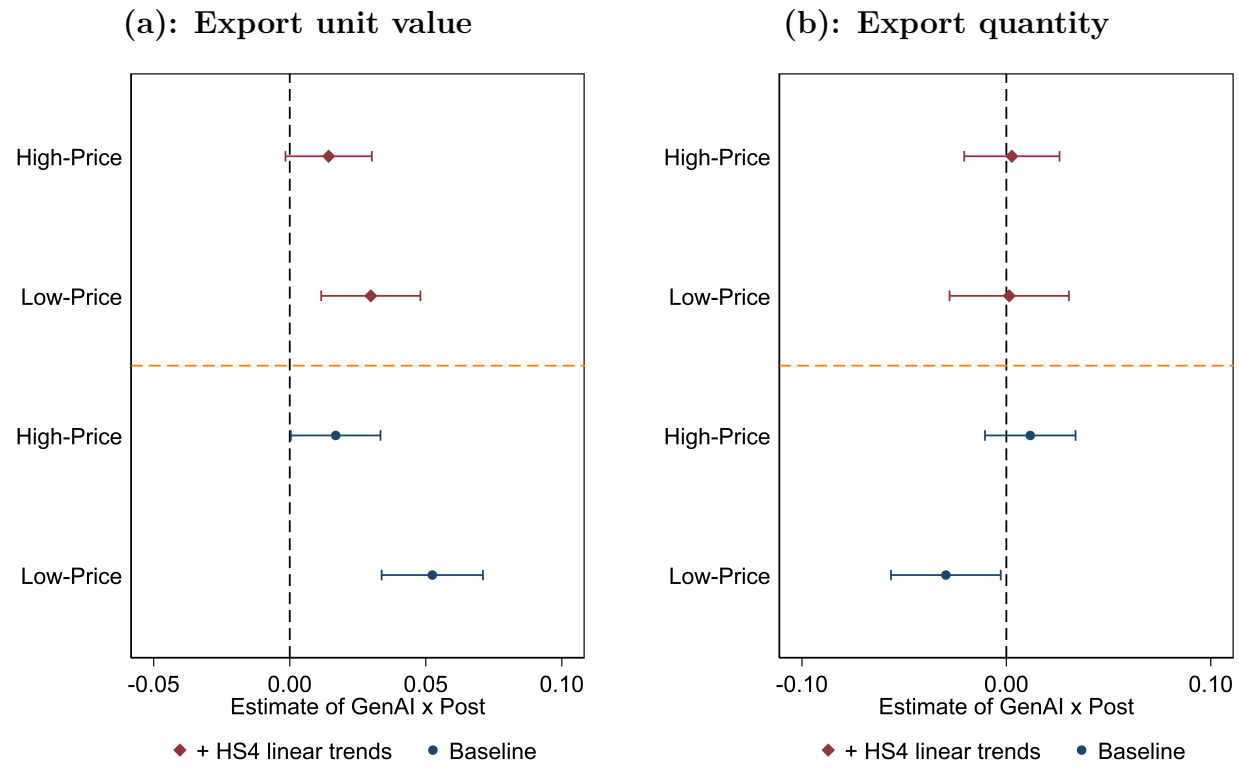
In sum, export unit values rise more among products with lower pre-shock unit values across both specifications, while the corresponding quantity decline is present in the baseline specification but is absorbed by HS4-specific linear month trends. These patterns are consistent with the cutoff-weighted comparative statics of Proposition 1 when the foreign-adoption selection channel is stronger, with the unit-value dimension surviving finer trend controls more robustly than the quantity dimension.

Table 4: Effects of GenAI Exposure by Pre-Period Product Unit Value within HS4

	(1)	(2)	(3)	(4)
	Mean Ln Aggregate Export Unit Value		Mean Ln Export Quantity	
	Low unit value	High unit value	Low unit value	High unit value
Panel A. HS10 Product FE + Month FE				
GenAI × Post	0.0524***	0.0169**	−0.0296**	0.0117
	(0.0095)	(0.0084)	(0.0137)	(0.0113)
Observations	165,975	193,239	165,975	193,239
R^2	0.9646	0.9703	0.9182	0.9245
Empirical p (low = high)	0.000		0.000	
Panel B. Adding HS4 Linear Trends				
GenAI × Post	0.0298***	0.0143*	0.0014	0.0027
	(0.0093)	(0.0081)	(0.0149)	(0.0119)
Observations	165,975	193,239	165,975	193,239
R^2	0.9663	0.9714	0.9207	0.9265
Empirical p (low = high)	0.020		0.910	
HS10 product FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Clustered SE (HS10)	Yes	Yes	Yes	Yes

Notes. The table reports difference-in-differences estimates of $\text{GenAI}_p \times \text{Post}_t$ from the HS10 product–month panel (2020m1–2025m12), estimated separately for low- and high-unit-value product subsamples within each HS4 category. Products are divided at the HS4-specific median of average aggregate export unit value in 2022m1–2022m11. Columns (1)–(2) use aggregate export unit value; columns (3)–(4) use mean export quantity. Panel A uses the baseline specification with HS10 product and year–month fixed effects; Panel B adds HS4-specific linear month trends. GenAI_p and Post_t are defined in Section 4. All specifications cluster standard errors at the HS10 product level. Empirical p-values test equality of the low- and high-unit-value coefficients from 100 random reassignments of the unit-value labels holding group sizes fixed. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Figure 4: Heterogeneous effects of GenAI exposure on aggregate export unit value and mean export quantity, by pre-release product unit value



Notes. The figure displays estimates of GenAI exposure on aggregate export unit value and mean export quantity at the product level from 2020m1 to 2025m12, with products divided into low-unit-value and high-unit-value groups at the HS4-specific median of average aggregate export unit value across HS10 products within the same HS4 category during 2022m1–2022m11. Markers plot coefficients on $\text{GenAI}_p \times \text{Post}_t$ for aggregate export unit value in panel (a) and mean export quantity in panel (b). Each panel stacks two specification blocks separated by a dashed orange horizontal rule: the upper block (maroon diamond markers) adds HS4-specific linear month trends; the lower block (navy circle markers) is the baseline with HS10 product and month fixed effects. Horizontal ranges show 95% confidence intervals based on standard errors clustered at the HS10 product level.

5.5 Heterogeneous Export Responses by Destination AI Intensity

The model in Section 3 distinguishes two AI channels that move U.S. export unit values and quantities to foreign destinations in opposite directions. Destination-side AI adoption increases foreign competitive pressure on marginal U.S. varieties, while U.S.-side AI adoption raises U.S. productivity. These channels imply different predictions for how the product-destination unit-value and quantity responses vary with destination AI intensity. We measure destination AI intensity using two complementary pre-release OECD/LinkedIn

indicators: the AI hiring rate, which captures destination-country demand for AI labor, and AI talent concentration, which captures the depth of the country's AI-skilled workforce. We split destinations at the cross-country median of each indicator and re-estimate the baseline difference-in-differences specification within each destination subsample.

Table 5 reports the estimates, and Figure 5 plots the same coefficients with 95% confidence intervals across the hiring split (panels a–b) and the talent split (panels c–d). Panel A splits destinations by the AI hiring rate, while Panel B splits destinations by AI talent concentration. We first examine export unit values. Across both splits, the unit-value response is small and statistically insignificant in low-AI destinations, but positive in high-AI destinations. Under the hiring-rate split, the high-AI coefficient is 0.0151 in the baseline specification and rises to 0.0303 after allowing for HS4-specific linear month trends; both estimates are statistically significant. Under the talent-concentration split, the high-AI coefficient remains positive in the baseline specification, at 0.0114, although it is not statistically significant. After adding HS4-specific trends, the coefficient increases to 0.0263 and becomes statistically significant. Randomization tests reject equality of the low- and high-AI unit-value coefficients in both specifications for both destination-intensity measures.

We next examine export quantities. In contrast to unit values, the quantity response is concentrated in low-AI destinations. The baseline coefficient is 0.0432 under the hiring-rate split and 0.0399 under the talent-concentration split, and both estimates are statistically significant. After adding HS4-specific trends, the coefficients remain positive, at 0.0222 and 0.0206, respectively. However, in high-AI destinations, the quantity coefficient is small and statistically insignificant across all specifications. Randomization tests reject equality of the low- and high-AI quantity coefficients across both specifications and both destination-intensity measures.

This pattern is consistent with destination-side AI adoption strengthening selection against marginal U.S. exporters in high-AI destinations. Export unit values rise primarily in high-AI destinations, where U.S. exporters face stronger AI-enabled foreign competition

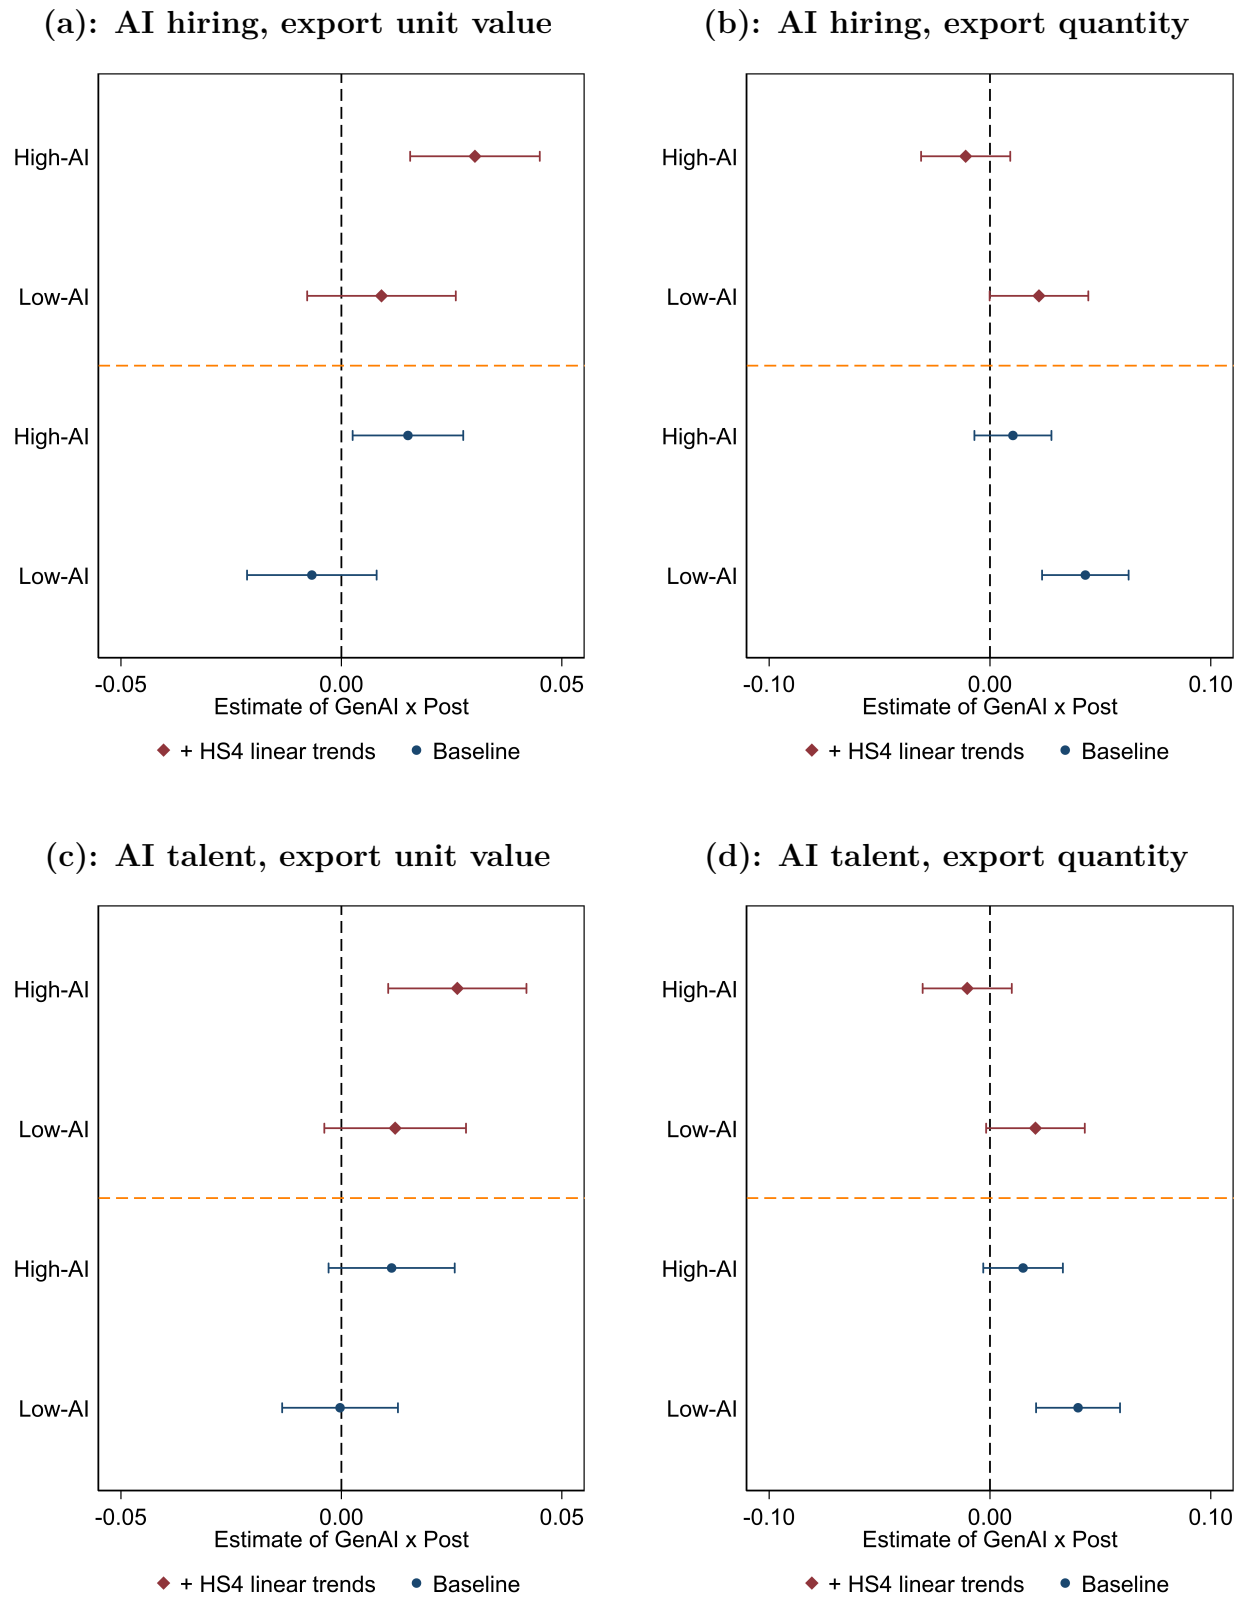
and quantities do not expand. Quantities increase primarily in low-AI destinations, where foreign-adoption pressure is weaker and the quantity-increasing effect of U.S. adoption can appear more clearly. Together with the pre-release unit-value heterogeneity results in subsection 5.4, the destination-level asymmetry supports the interpretation that the aggregate response reflects destination-side selection rather than a uniform expansion of U.S. export scale. Compared with the product unit-value split in subsection 5.4, the destination-AI split offers sharper resolution on the quantity margin, since the quantity asymmetry across destination groups survives HS4-specific linear month trends while the corresponding split across product unit values does not. The result directly supports the destination-side channel in Proposition 1.

Table 5: Effects of GenAI Exposure by Destination AI Intensity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Ln Export Unit Value				Ln Export Quantity			
	Baseline		With HS4 linear trend		Baseline		With HS4 linear trend	
	Low AI	High AI	Low AI	High AI	Low AI	High AI	Low AI	High AI
Panel A. OECD/LinkedIn AI hiring rate								
GenAI × Post	−0.0067	0.0151**	0.0091	0.0303***	0.0432***	0.0104	0.0222*	−0.0110
	(0.0075)	(0.0064)	(0.0086)	(0.0075)	(0.0100)	(0.0089)	(0.0114)	(0.0103)
Observations	1,933,425	2,049,312	1,933,425	2,049,312	1,936,783	2,052,598	1,936,783	2,052,598
R^2	0.7483	0.7775	0.7494	0.7785	0.6049	0.6823	0.6060	0.6832
Empirical p (low = high)	0.000		0.000		0.000		0.000	
Panel B. OECD/LinkedIn AI talent concentration								
GenAI × Post	−0.0003	0.0114	0.0122	0.0263***	0.0399***	0.0150	0.0206*	−0.0103
	(0.0067)	(0.0073)	(0.0082)	(0.0080)	(0.0097)	(0.0092)	(0.0114)	(0.0103)
Observations	1,888,235	2,111,840	1,888,235	2,111,840	1,891,359	2,115,443	1,891,359	2,115,443
R^2	0.7590	0.7687	0.7599	0.7698	0.6463	0.6630	0.6473	0.6639
Empirical p (low = high)	0.000		0.030		0.000		0.005	
HS10 product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × month trend	No	No	Yes	Yes	No	No	Yes	Yes
Clustered SE (HS10)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The table displays difference-in-differences estimates from the HS10 product–destination–month panel, estimated separately within low- and high-AI destination subsamples. Panel A partitions destinations at the cross-country median of pre-period OECD/LinkedIn AI hiring intensity; Panel B partitions destinations at the median of AI talent concentration. Columns (1)–(4) report Ln Export Unit Value and columns (5)–(8) report Ln Export Quantity; within each outcome block, columns (1)–(2) and (5)–(6) use the baseline specification while columns (3)–(4) and (7)–(8) add HS4-specific linear month trends, with each specification block further split between low- and high-AI destinations. GenAI_p and Post_t are defined in Section 4. All specifications include HS10 product, destination, and year–month fixed effects and cluster standard errors at the HS10 product level. Empirical p-values test equality of the low- and high-AI coefficients from 1,000 random reassignments of the destination labels holding group sizes fixed. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Figure 5: Heterogeneous effects of GenAI exposure on export unit value and export quantity, by destination AI intensity



Notes. The figure plots coefficients on $\text{GenAI}_p \times \text{Post}_t$ from product–destination regressions. Panels (a)–(b) split destinations by pre-period OECD/LinkedIn AI hiring rate, and panels (c)–(d) split them by AI talent concentration. Navy circles show the baseline estimates; maroon diamonds add HS4-specific linear month trends. Horizontal bars denote 95% confidence intervals with standard errors clustered at the HS10 product level.

5.6 Joint Heterogeneity by Pre-Release Product Unit Value and Destination AI Intensity

Table 6 crosses the two heterogeneity dimensions documented in subsections 5.4 and 5.5, pre-release product unit value and destination AI intensity, and re-estimates the baseline difference-in-differences specification within each group. Destinations are partitioned by hiring intensity in Panel A and by talent concentration in Panel B, and Figure 6 plots the same coefficients with 95% confidence intervals, making the joint pattern across product unit value and destination AI intensity visually transparent. The model implies that the two dimensions should compound: lower-unit-value products carry larger cutoff weights, and high-AI destinations impose stronger destination-side selection pressure.

The cross-tabulation shows that the two dimensions reinforce one another. Within low-unit-value products, the export unit-value response is positive in both destination groups but substantially larger and more precisely estimated in high-AI destinations. With HS4-specific linear month trends, the estimated coefficient for shipments to high-AI destinations is 0.0447 under the hiring-based split and 0.0446 under the talent-based split, both statistically significant at the 1% level, approximately double the corresponding coefficients for shipments to low-AI destinations, 0.0228 and 0.0224, which retain only marginal significance at the 10% level. Within high-unit-value products, the export unit-value response is statistically insignificant in every destination group once HS4-specific linear month trends are added, paralleling the muted high-unit-value coefficient in Table 4. The aggregate unit-value increase therefore concentrates among low-unit-value products shipped to high-AI destinations, precisely where the cutoff-weighted comparative statics of Proposition 1 should be strongest.

The quantity results show the other side of the same selection force. Quantity growth is concentrated among high-unit-value products shipped to low-AI destinations, where foreign-selection pressure is weaker and the positive quantity effect associated with U.S. adoption is less offset by foreign adoption. For high-unit-value products shipped to high-AI destinations,

the quantity response is much smaller, consistent with stronger destination-side AI adoption absorbing the expansionary force. For low-unit-value products, the quantity estimates are statistically weak across destination groups, but their signs follow the model's selection logic, with positive estimates in low-AI destinations and negative estimates in high-AI destinations. This sign pattern is consistent with low-unit-value products being more exposed to foreign selection pressure, especially where destination-side AI adoption is high.

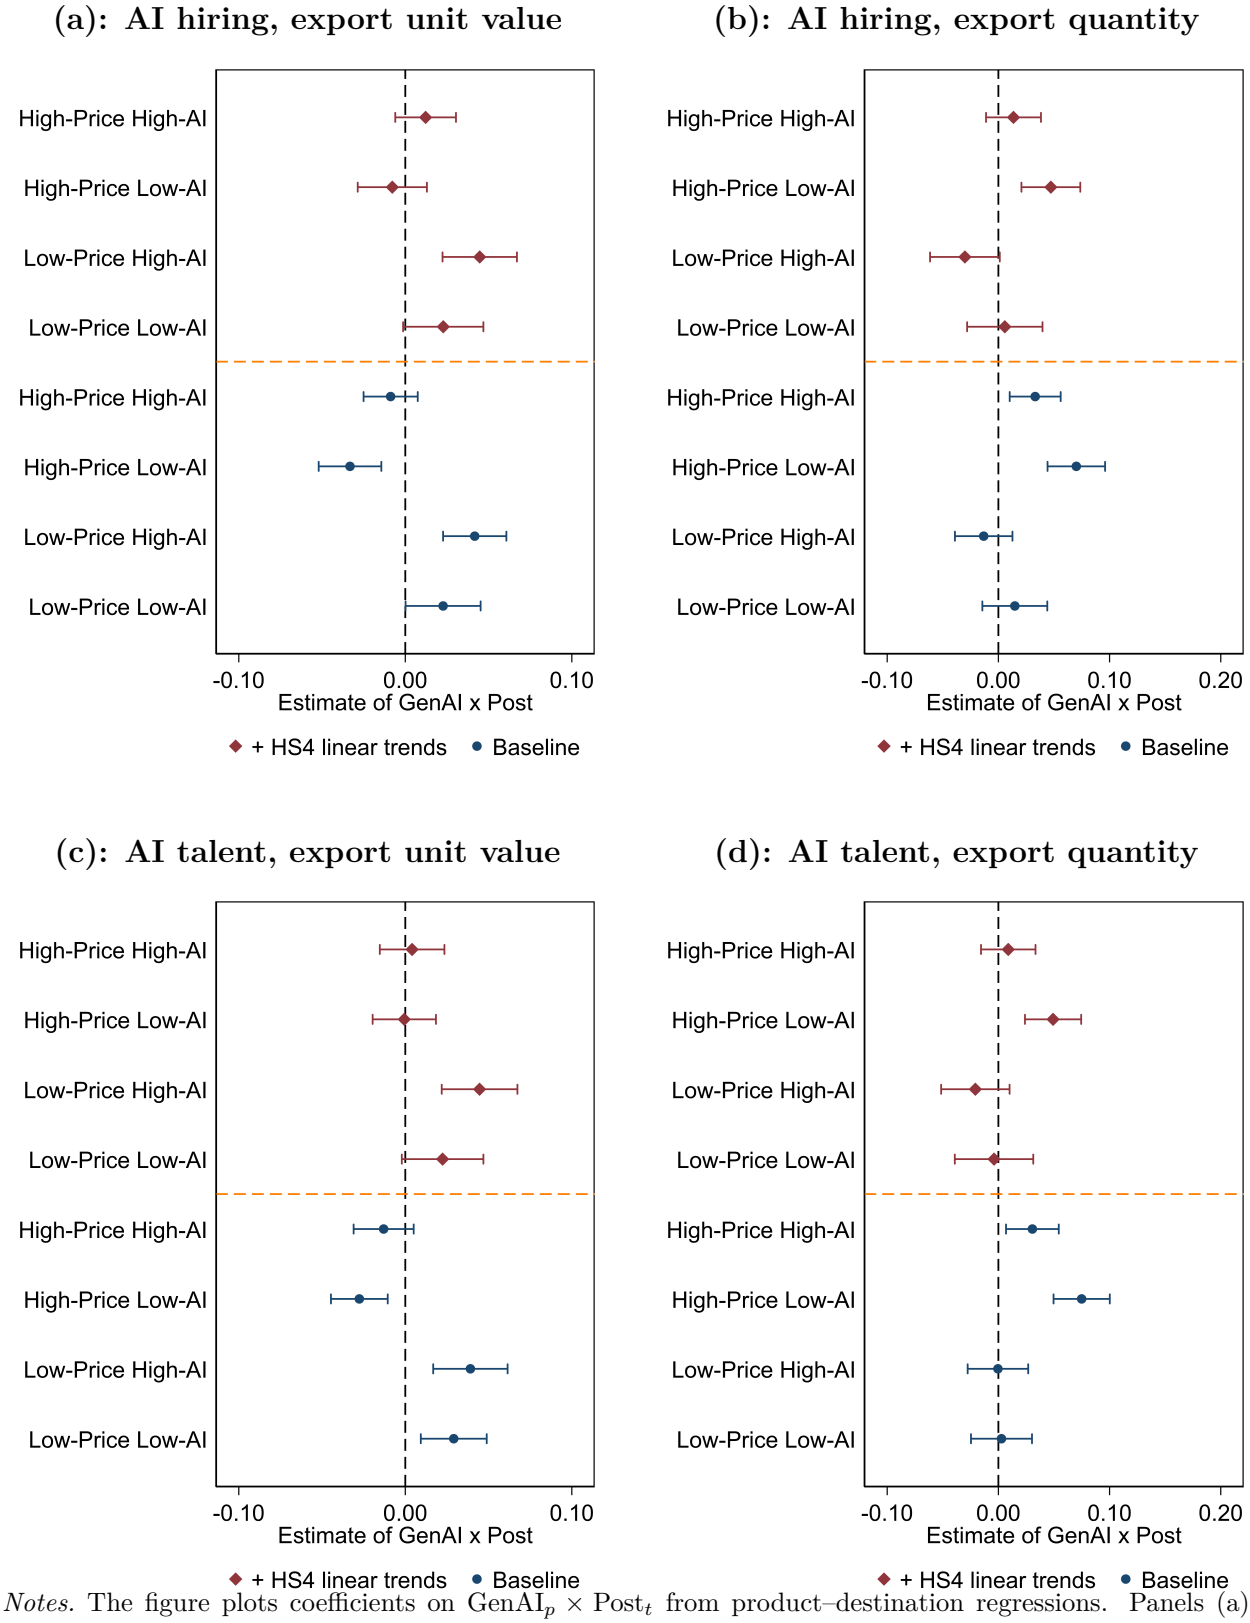
Taken together, the cross-tabulation supports the paper's main selection interpretation. The post-release rise in aggregate export unit values is concentrated where destination-side GenAI adoption should most strongly select lower-quality U.S. varieties out of export markets, while quantity growth appears mainly where that foreign-selection pressure is weakest.

Table 6: Effects of GenAI Exposure by Product Unit Value and Destination AI Intensity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Ln Export Unit Value				Ln Export Quantity			
	Baseline		With HS4 linear trend		Baseline		With HS4 linear trend	
	Low-AI	High-AI	Low-AI	High-AI	Low-AI	High-AI	Low-AI	High-AI
Panel A. AI Hiring Intensity								
Low-unit-value products								
GenAI × Post	0.0227**	0.0417***	0.0228*	0.0447***	0.0148	−0.0132	0.0058	−0.0301*
	(0.0115)	(0.0097)	(0.0123)	(0.0114)	(0.0149)	(0.0132)	(0.0173)	(0.0160)
Observations	878,794	927,219	878,794	927,219	878,796	927,219	878,796	927,219
High-unit-value products								
GenAI × Post	−0.0332***	−0.0088	−0.0078	0.0122	0.0701***	0.0331***	0.0472***	0.0136
	(0.0096)	(0.0083)	(0.0106)	(0.0093)	(0.0132)	(0.0117)	(0.0135)	(0.0126)
Observations	1,009,834	1,075,667	1,009,834	1,075,667	1,009,834	1,075,668	1,009,834	1,075,668
Panel B. AI Talent Concentration								
Low-unit-value products								
GenAI × Post	0.0291***	0.0391***	0.0224*	0.0446***	0.0028	−0.0004	−0.0039	−0.0207
	(0.0101)	(0.0114)	(0.0125)	(0.0116)	(0.0140)	(0.0139)	(0.0180)	(0.0157)
Observations	860,714	952,943	860,714	952,943	860,716	952,943	860,716	952,943
High-unit-value products								
GenAI × Post	−0.0276***	−0.0130	−0.0006	0.0041	0.0749***	0.0306**	0.0492***	0.0089
	(0.0087)	(0.0092)	(0.0097)	(0.0099)	(0.0129)	(0.0121)	(0.0129)	(0.0125)
Observations	984,292	1,110,363	984,292	1,110,363	984,292	1,110,364	984,292	1,110,364
HS10 product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × month trend	No	No	Yes	Yes	No	No	Yes	Yes
Clustered SE (HS10)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The table summarizes difference-in-differences estimates of GenAI exposure on export unit value and export quantity at the product-destination level from 2020m1 to 2025m12. Panel A partitions destinations at the cross-country median of pre-period OECD/LinkedIn AI hiring intensity, and Panel B partitions destinations at the cross-country median of pre-period AI talent concentration. Within each panel, products are classified at the HS4-specific median of average pre-period aggregate export unit value among HS10 products within the same HS4 category in 2022m1–2022m11, and the upper sub-block reports low-unit-value products while the lower sub-block reports high-unit-value products. Columns (1)–(4) report Ln Export Unit Value and columns (5)–(8) report Ln Export Quantity; within each outcome block, columns (1)–(2) and (5)–(6) use the baseline specification while columns (3)–(4) and (7)–(8) add HS4-specific linear month trends, with each specification block further split between low- and high-AI destinations. GenAI_p, Post_t, and the outcome definitions follow Equation (21) and Section 4. Continuous outcomes are winsorized at the 1st and 99th percentiles before taking logarithms. All specifications include HS10 product, destination, and year–month fixed effects and cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Figure 6: Heterogeneous effects of GenAI exposure on export unit value and export quantity, by pre-release product unit value and destination AI intensity



Notes. The figure plots coefficients on $\text{GenAI}_p \times \text{Post}_t$ from product–destination regressions. Panels (a) and (c) use export unit value as the outcome, while panels (b) and (d) use export quantity. Products are split into low- and high-unit-value groups using the HS4-specific median of average aggregate export unit value over 2022m1–2022m11. Destinations are split by pre-period OECD/LinkedIn AI hiring rate in panels (a)–(b) and AI talent concentration in panels (c)–(d). Navy circles report the baseline estimates; maroon diamonds add HS4-specific linear month trends. Horizontal bars show 95% confidence intervals with standard errors clustered at the HS10 product level.

5.7 Domestic Industry Organization

The export results above show that GenAI exposure changes which U.S. varieties remain active in foreign markets. Proposition 2 links the same global AI shock to the domestic mass of U.S. establishments producing in exposed industries. In the model, U.S. adoption and foreign adoption have opposing implications for establishment counts. U.S. adoption raises productivity in exposed products and can increase the expected value of serving export markets. Foreign adoption strengthens competition abroad in lower-quality market segments and reduces residual demand for U.S. exporters. The net effect on U.S. establishment counts is therefore ambiguous and depends on which force is stronger. Table 7 examines domestic industry organization after the public release of ChatGPT, including establishment counts, gross establishment entry, total employment, and average employment per establishment.

Table 7 reports Poisson estimates of GenAI exposure on four industry-level outcomes.[18] The estimates for gross establishment entry in column (2) are economically small and statistically insignificant. This entry result is the object most directly linked to Proposition 2: the model holds the exit rate fixed, so changes in establishment mass operate through the free-entry margin. The flat gross-entry response is consistent with the model's opposing-force logic, because the U.S. productivity force and the foreign selection force can offset one another on the entry margin. In column (1), the coefficient on establishment count is 0.0355 and statistically significant at the 5% level, equivalent to roughly 3.6% more establishments in more exposed industries after the public release of ChatGPT. Since gross entry does not rise, the positive establishment-count response is naturally interpreted through the exit margin: in the short run, exposed industries appear to retain more establishments, or lose fewer establishments, after the GenAI shock. Industry employment does not increase, and average employment per establishment falls. The establishment response therefore reflects a larger number of smaller operating units rather than an expansion of total industry scale.

[18] We follow Cohn, Liu, and Wardlaw (2022) in estimating count outcomes by Poisson pseudo-maximum-likelihood, which admits zero observations directly and is consistent under heteroskedasticity in the conditional mean, avoiding the bias of $\log(1+y)$ transformations.

Table 7: Effects of GenAI Exposure on Domestic Industry Organization

VARIABLES	(1) Establishments	(2) Entry	(3) Industry employment	(4) Avg. employment per establishment
GenAI × Post	0.0355**	0.0016	−0.0017	−0.0291***
	(0.0164)	(0.0085)	(0.0099)	(0.0102)
Observations	1,094	896	1,094	1,039
Pseudo R^2	0.9658	0.9929	0.9663	0.9962
Four-digit NAICS industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Clustered SE (four-digit NAICS)	Yes	Yes	Yes	Yes

Notes. The table displays Poisson regression estimates of GenAI-Exposure on several domestic industry-level variables from 2020 to 2025. Column (1) uses establishment count and column (3) uses industry employment, both summed from six-digit Quarterly Census of Employment and Wages observations within each four-digit industry; column (2) uses establishment entries from the Business Dynamics Statistics; column (4) uses average employment per establishment, estimated by Poisson with establishment count as an offset. The reported coefficient is $\text{GenAI}_i \times \text{Post}_t$. GenAI_i is pre-release GenAI industry exposure and Post_t equals one for years from 2023 onward, as in Section 4. Columns (1)–(3) apply the same Poisson estimator to count outcomes trimmed at the 1st and 99th percentiles before estimation. All specifications include four-digit industry and year fixed effects and cluster standard errors at the four-digit NAICS industry level. Standard errors are in parentheses; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

5.8 Industry-Destination Trade Structure

The model does not distinguish between arm's-length and related-party transactions. We therefore use the related-party data as an external diagnostic rather than as a direct test of Proposition 1. If the aggregate export unit-value increase mainly reflected transfer-pricing adjustments within multinational affiliate networks, the response should be more visible in related-party exports. If instead it reflects changes in market transactions with unaffiliated foreign buyers, the response should be more visible in arm's-length exports. Table 8 estimates the response separately by trade type using the Census industry–destination–year panel of arm's-length and related-party exports.

Column (1) shows that arm's-length exports increase significantly after the public release of ChatGPT, with an estimated coefficient of 0.1070. In contrast, Column (2) shows no statistically significant response in related-party exports. The increase in export val-

ues associated with GenAI exposure is therefore concentrated in arm's-length sales rather than in intra-firm transactions. This pattern is difficult to reconcile with a transfer-pricing explanation, under which related-party exports would be expected to respond more strongly.

Together with the establishment-count response documented in subsection 5.7, these results suggest that the aggregate response is unlikely to be driven by administrative price adjustments within multinational firms.

Table 8: Effects of GenAI Exposure on Arm's-Length Exports

VARIABLES	(1) Ln Arm's-Length Exports	(2) Ln Related-Party Exports
GenAI × Post	0.1070**	−0.0257
	(0.0523)	(0.0816)
Observations	58,767	58,767
Adjusted R^2	0.6249	0.6131
Four-digit NAICS industry FE	Yes	Yes
Destination FE	Yes	Yes
Year FE	Yes	Yes
Clustered SE (four-digit NAICS)	Yes	Yes

Notes. The table displays difference-in-differences estimates of GenAI-Exposure on Arm's-Length Exports and Related-Party Exports on industry level from 2020 to 2024. Column (1) uses arm's-length exports; column (2) uses related-party exports. Each outcome is the logarithm of one plus export value in the Census Related-Party Trade Data. The reported coefficient is $\text{GenAI}_i \times \text{Post}_t$. GenAI_i is pre-release GenAI industry exposure and Post_t equals one for post-release years, as in Section 4. All specifications include four-digit NAICS industry, destination, and year fixed effects and cluster standard errors at the four-digit NAICS industry level. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

6 Robustness Checks

6.1 Continuous-Treatment Robustness

Our baseline specification in Equation (21) interacts a continuous product-level GenAI exposure measure with the post-release indicator. As Callaway, Goodman-Bacon, and Sant'Anna (2024) note, when treatment intensity varies continuously, the standard two-way fixed-effects estimator aggregates effects across treatment dosages and can be difficult to interpret if the

underlying dose-response relationship is nonlinear. We therefore examine whether the baseline finding is robust to alternative ways of modeling treatment intensity.

First, we replace the continuous GenAI exposure measure with a high-exposure indicator GenAIHigh_p that equals one for HS10 products whose pre-shock GenAI exposure exceeds the specified percentile cutoff and equals zero otherwise. The bottom decile, and alternatively the bottom 15, 20, and 25 percent, of the pre-shock exposure distribution serves as the reference group. This approach replaces the linear continuous-interaction term with a comparison between high- and low-exposure products. Appendix Table A.1 shows that findings are unchanged in every cutoff–specification combination. The estimates therefore do not depend on the particular percentile chosen to demarcate the low-exposure reference group.

Second, we examine whether the baseline result depends on imposing a linear relationship between GenAI exposure and the outcome. We discretize the continuous exposure measure into ordered bins and estimate bin-specific difference-in-differences coefficients, using the lowest-exposure bin as the omitted category. This specification allows the post-shock response to vary flexibly across the exposure distribution. Appendix Figures A.1(a)–(d) plot the estimated coefficients from specifications with five to eight exposure bins. Across all binning choices, the coefficients increase monotonically with exposure. The estimates become larger in the higher-exposure bins, consistent with a positive relationship between GenAI exposure and the post-shock increase in export unit values.

Third, we implement the nonparametric continuous-treatment difference-in-differences estimator of Callaway, Goodman-Bacon, and Sant'Anna (2024), which allows the dose-response function to be flexible. Because the GenAI exposure measure has no natural zero-dose benchmark, we use HS10 products in the bottom decile, and alternatively in the bottom 15, 20, and 25 percent, of the pre-shock exposure distribution as the least-exposed comparison group. The estimated average treatment effects are positive, statistically significant, and consistent with the baseline (Appendix Table A.2); the dose-response curve is positive across the support of the exposure distribution (Appendix Figure A.2); and the cor-

responding event-study coefficients display flat pre-shock dynamics and positive post-shock dynamics (Appendix Figure A.3), mirroring the baseline event study in Figure 3.

Fourth, we address measurement error in the exposure assignment by implementing the fuzzy difference-in-differences estimator of de Chaisemartin and D'Haultfoeuille (2018), which recovers the local average treatment effect on the population of switchers under imperfect treatment assignment. Switchers are HS10 products that move above the indicated lower-tail cutoff of the exposure distribution; stayers remain below. Appendix Table A.3 shows that the estimated local average treatment effects are consistent with the baseline.

Overall, the estimated effect of GenAI exposure on aggregate export unit values is stable across alternative ways of modeling exposure intensity, alternative low-exposure comparison groups, and corrections for measurement error in exposure assignment, indicating that the baseline finding is not driven by a particular functional-form assumption or comparison-group definition.

6.2 Alternative Exposure Sources

A related concern is whether the post-release increase in export unit values reflects image-generation capabilities rather than the language-, communication-, and information-processing capabilities targeted by GenAI_p. We construct an industry-level image-generation exposure measure from the AI Industry Exposure dataset of Felten, Raj, and Seamans (2021) using the same occupational–industry aggregation and HS10 concordance as GenAI_p, and we re-estimate the baseline specification with both interactions entering simultaneously. Appendix Table A.4 shows that the GenAI coefficient remains positive and statistically significant in both specifications, while the image-generation coefficient is close to zero and statistically insignificant. The post-release unit-value response therefore aligns with language-model exposure rather than image-generation exposure.

6.3 Placebo Tests

To assess whether the baseline estimate could plausibly arise by chance, we replace the actual product-level GenAI exposure with a placebo, interact it with the post-release indicator, and re-estimate the baseline specification 1,000 times under three randomization designs—permuting the actual exposure vector across HS10 products, drawing placebo exposure from $U(0,1)$, and drawing it from $U(0, \max_p \text{GenAI}_p)$. The placebo distributions are approximately centered at zero, the actual baseline estimate lies far in the right tail (Appendix Figure A.4), and the empirical p-values are small (Appendix Table A.5), indicating that the baseline coefficient is unlikely to be generated by random assignment of GenAI exposure across products.

6.4 Industry-Level Clustering

GenAI exposure is assigned to HS10 products through their corresponding six-digit NAICS industry, so regression residuals may co-move across products that share an industry. To allow for this within-industry correlation, we re-cluster standard errors at the four-digit NAICS level. Appendix Table A.6 shows that the estimates are unchanged.

7 Quantitative Model Analysis

The reduced-form estimates identify the direction and heterogeneity of the post-ChatGPT responses, but they do not directly map those responses into the structural selection channel in the model. This section describes the calibration strategy that takes the model of Section 3 to the data, reports the estimated AI-shock parameters, and uses the estimated model to decompose price selection and compute welfare and aggregate price effects.

7.1 Calibration

The structural parameters fall into four groups, calibrated in order: preset constants and normalizations, the iceberg trade cost, time-invariant fundamentals recovered from the baseline equilibrium given η, and four AI-shock elasticity parameters $(\eta, \alpha, \rho, \nu)$ estimated by simulated method of moments. A separate block lists the observed AI-shock inputs that drive the post-ChatGPT simulation. Table 9 summarizes the identification of each object, and the SMM algorithm, data-cleaning pipeline, and detailed closed-form solution are deferred to Appendix D. The simulated equilibria for candidate AI paths use the solving algebra in Appendix C.

Preset constants and normalizations: $\{\sigma,\ k_j/\eta,\ \delta_{\text{mo}},\ \underline{a}_j\}$. The CES substitution elasticity $\sigma = 4$ is the median of the distribution across products estimated by Broda and Weinstein (2006). The trade elasticity is the iceberg elasticity of bilateral revenue implied by the model: the baseline revenue equation derived in Section 3.2 satisfies $R^x_{dj,0} \propto \tau_{dj}^{-k_j/\eta}$, so k_j/η is the trade elasticity. We borrow $k_j/\eta = 4.25$ from Simonovska and Waugh (2014); this single restriction pins the ratio and implies $k_j = 4.25\,\eta$ once η is estimated. The stationary monthly establishment exit rate $\delta_{\text{mo}} = 0.0087$ is the U.S. manufacturing annual establishment exit rate of approximately 0.10 from the U.S. Census Bureau Business Dynamics Statistics,[19] converted to monthly frequency via $\delta_{\text{mo}} = 1-(1-0.10)^{1/12}$. The Pareto lower bound $\underline{a}_j = 1$ is a scale normalization on within-industry firm-type units.

Iceberg trade cost: $\{\tau_{dj}\}$. We estimate τ_{dj} by decomposing it as $\tau_{dj} = \tau_d\,\tau_j$ and recovering each component separately. The destination component τ_d is identified from the pre-AI closed-form export revenue relationship $R^x_{dj,0} \propto \tau_{dj}^{-k_j/\eta}$. Substituting the closed-form foreign zero-profit cutoff into the baseline revenue identity, summing over products, and taking logs delivers the gravity reduced form

$$\log X_{d,0} \;=\; \alpha_0 \;-\; \tfrac{k_j}{\eta}\log\tau_d \;+\; \Phi(Y_d, \mathbf{X}_d) \;+\; u_d, \qquad X_{d,0} \equiv \textstyle\sum_j R^x_{dj,0}, \tag{23}$$

[19] U.S. Census Bureau, Business Dynamics Statistics (BDS), `www.census.gov/programs-surveys/bds.html`.

in which $\log \tau_d$ is proxied by bilateral resistance variables from the CEPII Gravity database[20] and $\Phi(Y_d, \mathbf{X}_d)$ absorbs destination size and other demand shifters. The full derivation of (23) and the gravity specification are in Appendix D.2. The fitted iceberg coefficient is rescaled by $-1/(k_j/\eta)$ and the level is normalized so that $\hat{\tau}_{\text{CAN}} = 1$. The industry multiplier $\hat{\tau}_j$ is estimated from the ASM revenue–labor share ratio. Let r_j and ℓ_j be industry j's normalized revenue and labor-cost shares, and let ϕ_j^H and ϕ_{dj}^x be the home and export revenue shares within the industry. The calibration first inverts $\hat{\bar{\tau}}_j = \sigma(r_j/\ell_j)/\{\mu[\sigma - (r_j/\ell_j)]\}$, then computes $\bar{\tau}_j^d = (\phi_j^H + \sum_d \phi_{dj}^x/\hat{\tau}_d)^{-1}$, and sets $\hat{\tau}_j = \hat{\bar{\tau}}_j/\bar{\tau}_j^d$. The full bilateral iceberg used downstream is $\hat{\tau}_{dj} = \hat{\tau}_d\,\hat{\tau}_j$. Appendix D.2 reports the $\hat{\tau}_j$ estimation and data construction.

Time-invariant fundamentals given η: $\{\hat{A}_j, \widehat{\Gamma}_{dj,0}, \hat{f}_{dj}, \hat{f}_j^H, \hat{\beta}_j, \hat{f}_j^E\}$. Given a trial value of η and the constants above, the time-invariant fundamentals of Section 3.2 are recovered in closed form from observed pre-period objects. The industry productivity scale $\hat{A}_j$ is identified from the product–destination export unit value $p_{dj,0}^x$ together with the industry-level successful-entry probability s_j^H from the U.S. Census Business Formation Statistics[21] and the exporter share π_j from the 2022 Census exporter survey,[22] which together pin the home and foreign cutoffs and hence the absolute scale of the Pareto support. Conditional on $\hat{A}_j$, the foreign residual-demand shifter $\widehat{\Gamma}_{dj,0}$ follows from the baseline export revenue identity; the fixed market-access costs $\hat{f}_{dj}$ and $\hat{f}_j^H$ follow from the corresponding zero-profit cutoff conditions; the Cobb–Douglas expenditure share is $\hat{\beta}_j = R_{j,0}^H/L_t$, with home revenue from the Annual Survey of Manufactures[23] and aggregate income from the U.S. Bureau of Economic Analysis monthly personal-income series;[24] and the sunk entry cost $\hat{f}_j^E$ follows from the free-entry condition. The detailed solution of the baseline system and the data

[20] Conte, M., P. Cotterlaz, and T. Mayer (2022). *The CEPII Gravity Database*, CEPII Working Paper No. 2022-05. `www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=8`.

[21] U.S. Census Bureau, Business Formation Statistics (BFS), `www.census.gov/econ/bfs/`.

[22] U.S. Census Bureau, "Profile of U.S. Importing and Exporting Companies, 2022," `www.census.gov/foreign-trade/Press-Release/edb/`.

[23] U.S. Census Bureau, Annual Survey of Manufactures (ASM), `www.census.gov/programs-surveys/asm.html`.

[24] U.S. Bureau of Economic Analysis, Personal Income series `PI`, retrieved through FRED, `fred.stlouisfed.org/series/PI`.

construction steps used on each input are reported in Appendix D.3.

AI-shock inputs: $\{s_j,\, M_{Ut},\, M_{dt}\}$. The post-ChatGPT shock is driven by three observed inputs that are the empirical counterparts of the model's exposure primitives and are held fixed while the elasticities are estimated. Product exposure s_j is the product-level GenAI exposure used in the reduced-form analysis, constructed by mapping the industry AI-exposure measure of Felten, Raj, and Seamans (2021) to HS10 products. U.S. adoption M_{Ut} is the post-release indicator, equal to zero before the ChatGPT release and one after. Destination adoption M_{dt} is the destination's pre-release AI intensity, measured by the OECD/LinkedIn AI hiring rate. These inputs enter the AI shock $\psi_{cjt}(a) = M_{ct}\, s_j\, \zeta_\eta(\xi(a))$ of Section 3; the elasticities estimated next govern how strongly that shock translates into export outcomes.

Elasticity parameters: $\{\eta,\, \alpha,\, \rho,\, \nu\}$. The four AI-shock parameters $\theta = (\eta, \alpha, \rho, \nu)$ are estimated by simulated method of moments using four reduced-form coefficients from Section 5. The first moment is the aggregate HS10 unit-value response: after exports are aggregated across destinations within each product, it measures whether products with greater GenAI exposure experience larger post-ChatGPT changes in unit values. The second moment is the unit-value response in the part of the product–destination distribution most exposed to selection: products with below-median pre-shock unit values shipped to destinations with above-median pre-shock AI hiring. The third moment is the corresponding quantity response for the same product–destination markets. The fourth moment is the aggregate HS10 quantity response, which measures whether more exposed products experience larger post-ChatGPT changes in total export quantities after aggregating over destinations. These four moments discipline the average price response, the selection component of prices, the selection component of quantities, and the aggregate quantity response.

For every candidate θ, the closed-form pre-shock calibration above is re-evaluated at the trial η so that all primitives adapt to the proposed η. The post-shock equilibrium is then resolved under the post-period values of the AI-shock inputs, and the model-implied $\hat{\beta}_{\text{model}}(\theta)$ is computed from the simulated two-period first-difference cross-section. The estimator min-

imizes the SE-weighted distance between empirical and simulated coefficients,

$$\hat{\theta} = \arg\min_{\theta} \sum_{m} \frac{\left(\hat{\beta}_{m,\text{model}}(\theta) - \hat{\beta}_{m,\text{emp}}\right)^2}{\widehat{\text{SE}}^{\,2}(\hat{\beta}_{m,\text{emp}})},$$

subject to the regularity bounds $\eta > 0$, $\alpha, \nu \geq 0$, and $\rho > 0$. This procedure gives $(\hat{\eta}, \hat{\alpha}, \hat{\rho}, \hat{\nu}) = (0.0567, 0.0041, 0.1254, 0.0062)$. Table 9 summarizes the role of each parameter. Appendix D.4 reports the nested search algorithm, targeted coefficients, and inference procedure.

Table 9: Parameters and identification

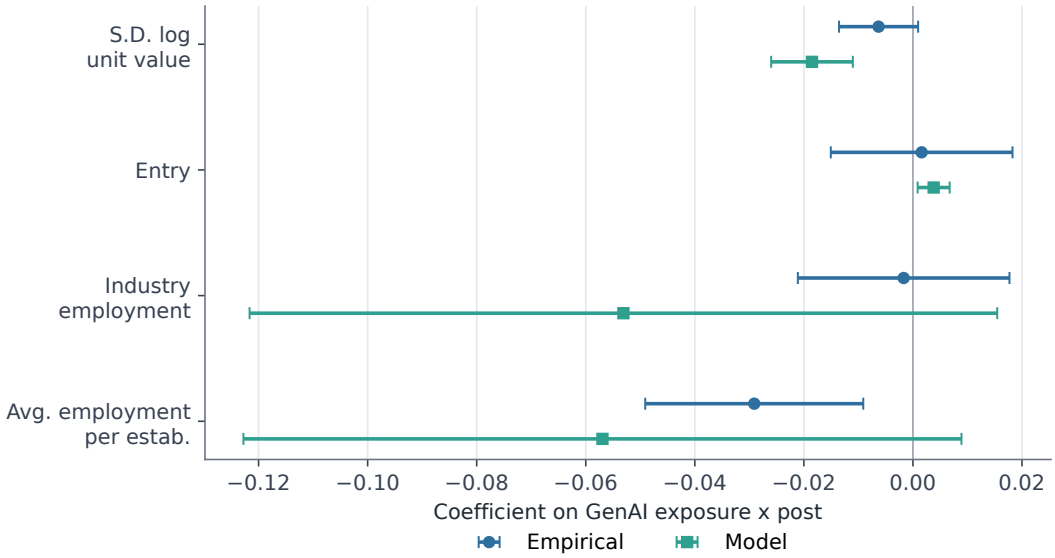
Parameter	Role	Identification	Estimate
Preset constants and normalizations.			
σ	CES substitution elasticity	Broda and Weinstein (2006)	4
k_j/η	Trade elasticity	Simonovska and Waugh (2014)	4.25
δ_{mo}	Monthly establishment exit rate	Census BDS	0.0087
$\underline{a}_j$	Pareto lower bound	Normalization	1
Iceberg trade cost.			
$\hat{\tau}_{dj}$	Iceberg trade cost	Gravity and revenue–labor inversion (Appendix D.2)	—
Time-invariant fundamentals.			
$\hat{A}_j$	Industry productivity scale	Model inversion (Appendix D.3)	—
$\hat{\Gamma}_{dj,0}$	Foreign residual-demand shifter		—
$\hat{f}_{dj},\ \hat{f}_j^H$	Fixed market-access costs		—
$\hat{\beta}_j$	Cobb–Douglas expenditure share		—
$\hat{f}_j^E$	Sunk entry cost		—
AI-shock inputs.			
s_j	Product GenAI exposure	Felten, Raj, and Seamans (2021)	—
M_{Ut}	U.S. AI adoption	Post-ChatGPT indicator, normalized	1
M_{dt}	Destination AI adoption	OECD/LinkedIn AI hiring rate	—
Estimated elasticities.			
$\hat{\eta}$	Quality elasticity in $\xi(a) = a^{1+\eta}$	4-moment SMM (Appendix D.4)	0.0567
$\hat{\alpha}$	Reduction in foreign demand from destination AI		0.0041
$\hat{\rho}$	How AI effect varies across the quality ladder		0.1254
$\hat{\nu}$	Rise in U.S. productivity from U.S. AI		0.0062

7.2 Model Validation

The SMM exercise targets only export-price and export-quantity moments: the aggregate HS10 unit-value response, the unit-value response among low-price products shipped to high-AI destinations, the corresponding quantity response, and the aggregate HS10 quan-

tity response. No domestic-market moment enters the objective function. Domestic organization outcomes therefore serve as model-validation outcomes. For each outcome, the model-predicted variable is constructed using the same definition as in the empirical analysis, and the corresponding first-difference regression is estimated on model-generated pre- and post-shock outcomes.

Figure 7: Model Validation: Regressions Using Model-Predicted Outcomes



Notes. The figure compares empirical coefficients with coefficients estimated from model-predicted outcomes at the preferred SMM parameters. Horizontal ranges are 95% confidence intervals. Empirical intervals use the standard errors from the reduced-form regressions; model intervals condition on the SMM parameters and use the same first-difference validation regressions applied to the model-generated data.

The validation exercise is informative because the domestic-market coefficients are not used to estimate $(\eta, \alpha, \rho, \nu)$. The model nevertheless generates entry and labor-market responses that are statistically compatible with the reduced-form estimates: exposed industries do not exhibit a clear total-employment expansion, while average employment per establishment falls.

7.3 Mechanism Decomposition and Welfare

Table 10 uses the estimated model to report two price-decomposition exercises. The first separates the source of the GenAI shock: we re-solve the model with only U.S. adoption, with

only destination-market adoption, and with both shocks. The second removes the quality-biased component of the shock by setting $\rho = 0$, so that AI exposure is common across qualities within a product. Each entry is a model-implied regression coefficient. For each exercise, we solve the pre- and post-shock equilibria, construct simulated trade outcomes, and regress the simulated post-minus-pre price change on product-level GenAI exposure using the same outcomes as in the reduced-form analysis.

Table 10: Model-Implied Regression Coefficients for Price Decomposition

Exercise	Aggregate HS10 unit value	Low-price, high-AI unit value	S.D. log unit value
Shock-source decomposition			
U.S. AI only	-0.0276^{***}	-0.0194^{***}	0.0043^{***}
Destination AI only	0.0490^{**}	0.0867^{***}	-0.0241^{***}
Full shock	0.0239	0.0698^{***}	-0.0185^{***}
Quality-biased mechanism			
$\rho = 0$	0.1745^{***}	0.4558^{***}	0.1205^{***}

Notes. Each entry is the coefficient from a reduced-form regression run on model-generated two-period data. For each row, the model is solved at the pre-shock allocation and under the shock or restriction listed in the row; the simulated post-minus-pre change in the relevant price outcome is then regressed on product GenAI exposure. A coefficient of b means that a product with one unit higher GenAI exposure has a b log-point larger simulated change in that outcome. The low-price/high-destination-AI column uses product-destination observations for products with below-median pre-shock unit values shipped to destinations with above-median pre-shock AI hiring. The dispersion outcome is the cross-destination standard deviation of log unit values within a product-month. The $\rho = 0$ row removes the quality gradient in AI exposure while retaining the estimated price-quality mapping. Stars follow the empirical tables: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The full-shock row is the model counterpart of the empirical price response. The model generates a heterogeneous price response: aggregate HS10 unit values rise on average, the response is larger for low-price products shipped to high-AI destinations, and cross-destination unit-value dispersion falls. The source decomposition explains this pattern. If only U.S. AI adoption is switched on, the signs move in the opposite direction: aggregate unit values and low-price, high-AI unit values fall, while dispersion rises. This is the natural implication of a U.S.-side productivity shock that is stronger for lower-quality varieties. Exports of lower-quality, lower-price varieties increase more, so average unit values fall, with the largest decline in the low-price, high-AI-destination sample. Destination AI works in the opposite

direction. It reduces foreign demand for U.S. varieties, with a larger effect on lower-quality varieties. Low-price U.S. exports therefore account for a smaller share of observed trade, average unit values rise, and dispersion falls. The full-shock outcome is therefore driven mainly by the destination-side force, with the U.S.-side force offsetting part of the price increase.

The quality-mechanism row sets $\rho = 0$, which removes the heterogeneous effect of AI across qualities within a product. Once this quality-biased force is removed, the model loses the empirical pattern. All three coefficients increase: the aggregate unit-value coefficient rises from 0.024 to 0.175, the low-price/high-AI coefficient rises from 0.070 to 0.456, and the dispersion coefficient changes from a decline (-0.019) to an increase (0.121). The model then predicts broad price increases and rising cross-destination price dispersion, rather than the observed combination of a modest aggregate price increase, a stronger low-price/high-AI response, and falling dispersion. This contrast shows that the quality-biased AI shock is central to the quantitative mechanism.

We use the estimated model to quantify the aggregate price and welfare effects of GenAI adoption.[25] Table 11 shows that welfare is mainly governed by the balance between domestic productivity gains and foreign competitive pressure. U.S. AI alone raises welfare by 1.19 percent because it makes U.S. establishments more productive and raises profits. Foreign AI alone lowers welfare by 0.54 percent because it strengthens foreign competitors and reduces the profitability of U.S. export sales. Under the full GenAI shock, these two forces operate together. The domestic productivity gain is larger than the foreign competitive loss for welfare, so the net welfare effect is positive, at 0.53 percent.

The price counterfactuals show how the quality-biased channel affects aggregate export unit values. U.S. AI alone lowers aggregate export unit values by 8.15 percent because productivity gains are stronger for lower-quality varieties, which expands lower-price U.S. exports and pulls down the average export price. Foreign AI alone raises aggregate export unit values by 17.28 percent because destination-side adoption hits lower-quality varieties

[25] In the model counterfactuals, aggregate export unit value is computed as total model export value divided by total model export quantity. The table reports the percent change in this aggregate ratio.

more strongly, reducing the export share of lower-price U.S. varieties and pushing up the average price of observed exports. The full GenAI shock combines these opposing price forces. The foreign-selection force is stronger, so aggregate export unit values rise by 7.54 percent. Setting $\rho = 0$ lowers the aggregate price effect to 1.78 percent, showing that quality-biased exposure is central to the aggregate price response.

Table 11: Welfare and Price Counterfactuals

Counterfactual	Welfare	Export unit value
U.S. AI only	1.19	-8.15
Foreign AI only	-0.54	17.28
Full shock	0.53	7.54
$\rho = 0$	0.31	1.78

Notes. Entries are percent changes from model levels. Each counterfactual compares its own pre-shock and post-shock model outcomes. Export unit value is total model export value divided by total model export quantity. The $\rho = 0$ row sets $\rho = 0$ only inside the AI-shock terms.

8 Conclusion

This paper studies how GenAI affects U.S. export prices and identifies the channel through which these effects operate. The channel comes from global, quality-biased adoption. In the model, U.S. establishments differ in cost and quality, so export unit values are composition outcomes shaped by selection into foreign destinations. GenAI enters through two adoption margins, and in both places the shock is stronger for lower-quality varieties. U.S. adoption raises exporter productivity and lowers the export threshold faced by U.S. firms. Destination adoption strengthens lower-quality foreign producers, raises the threshold faced by U.S. exporters abroad, and selects lower-quality U.S. varieties out of foreign markets. These two forces have opposite implications for prices and quantities, so GenAI can change the composition of trade without producing broad export growth.

The evidence supports this mechanism. After the release of ChatGPT, products with

greater GenAI exposure record higher export unit values, while average quantities remain broadly unchanged. The price response is strongest among initially low-unit-value products and in destinations with higher pre-release AI intensity. In the joint split, the largest unit-value increase appears for initially low-unit-value products shipped to high-AI destinations, precisely where the model predicts the strongest destination-side selection. Cross-destination unit-value dispersion falls, and the quantity response is concentrated in low-AI destinations, where foreign selection pressure is weaker. Domestic gross entry changes little, which is consistent with offsetting U.S. productivity and foreign selection forces on the entry margin, while the establishment count rises in a way that points to short-run continuation among existing establishments. The estimated model attributes the export-price increase mainly to destination-side selection, shows that the quality gradient in AI exposure is necessary to match the low-price high-AI response, and implies a positive net U.S. welfare effect of 0.53 percent.

The analysis also leaves scope for future work. Customs data measure unit values rather than product quality directly, and the post-ChatGPT window captures early adjustment rather than long-run dynamics. Future research combining firm-level AI adoption data with customs transactions and buyer-supplier links could study how generative AI reshapes exporter selection, firm organization, and the distribution of gains from trade over longer horizons.

References

Ali, Hassnian, Atta ul Mustafa, and Ahmet Faruk Aysan, 2025, Global adoption of generative AI: What matters most?, *Journal of Economy and Technology* 3, 166–176.

Allen, Treb, 2014, Information frictions in trade, *Econometrica* 82, 2041–2083.

Amiti, Mary, and Amit K. Khandelwal, 2013, Import competition and quality upgrading, *Review of Economics and Statistics* 95, 476–490.

Antoniades, Alexis, Chuan He, Zheming Liang, and Mingzhi Jimmy Xu, 2026, Learning by AI: Market intelligence and exporting, CESifo Working Paper 12456 CESifo.

Atkin, David, Amit K. Khandelwal, and Adam Osman, 2017, Exporting and firm performance: Evidence from a randomized experiment, *Quarterly Journal of Economics* 132, 551–615.

Autor, David H., David Dorn, and Gordon H. Hanson, 2013, The China syndrome: Local labor market effects of import competition in the United States, *American Economic Review* 103, 2121–2168.

Azoulay, Pierre, Joshua L. Krieger, and Abhishek Nagaraj, 2024, Old moats for new models: Openness, control, and competition in generative AI, *NBER Working Paper* pp. –.

Baldwin, Richard, and James Harrigan, 2011, Zeros, quality, and space: Trade theory and trade evidence, *American Economic Journal: Microeconomics* 3, 60–88.

Bastos, Paulo, Joana Silva, and Eric Verhoogen, 2018, Export destinations and input prices, *American Economic Review* 108, 353–392.

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott, 2006, Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of U.S. manufacturing plants, *Journal of International Economics* 68, 219–237.

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, 2011, Multiproduct firms and trade liberalization, *Quarterly Journal of Economics* 126, 1271–1318.

Bick, Alexander, Adam Blandin, and David J. Deming, 2024, The rapid adoption of generative AI, NBER Working Paper 32966 National Bureau of Economic Research.

Blum, Bernardo S., and Avi Goldfarb, 2006, Does the Internet defy the law of gravity?, *Journal of International Economics* 70, 384–405.

Bonfiglioli, Alessandra, Rosario Crinò, Mattia Filomena, and Gino Gancia, 2025, Comparative advantage in AI-intensive industries: Evidence from US imports, CESifo Working Paper 11642 CESifo.

Brambilla, Irene, Daniel Lederman, and Guido Porto, 2012, Exports, export destinations, and skills, *American Economic Review* 102, 3406–3438.

Broda, Christian, and David E. Weinstein, 2006, Globalization and the gains from variety, *Quarterly Journal of Economics* 121, 541–585.

Brynjolfsson, Erik, Xiang Hui, and Meng Liu, 2019, Does machine translation affect international trade? evidence from a large digital platform, *Management Science* 65, 5449–5460.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, 2025, Generative AI at work, *Quarterly Journal of Economics* 140, 889–942.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro, 2019, Trade and labor market dynamics: General equilibrium analysis of the China trade shock, *Econometrica* 87, 741–835.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, 2024, Difference-in-differences with a continuous treatment, NBER Working Paper 32117 National Bureau of Economic Research.

Chaney, Thomas, 2008, Distorted gravity: The intensive and extensive margins of international trade, *American Economic Review* 98, 1707–1721.

Cohn, Jonathan B., Zack Liu, and Malcolm I. Wardlaw, 2022, Count (and count-like) data in finance, *Journal of Financial Economics* 146, 529–551.

Crozet, Matthieu, Keith Head, and Thierry Mayer, 2012, Quality sorting and trade: Firm-level evidence for French Wine, *Review of Economic Studies* 79, 609–644.

de Chaisemartin, Clement, and Xavier D'Haultfoeuille, 2018, Fuzzy differences-in-differences, *Review of Economic Studies* 85, 999–1028.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik, 2016, Prices, markups, and trade reform, *Econometrica* 84, 445–510.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, 2026, Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality, *Organization Science* 37, 403–423.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang, 2023, Generative AI and firm values, NBER Working Paper 31222 National Bureau of Economic Research.

Fan, Haichao, Yao Amber Li, Sichuang Xu, and Stephen R. Yeaple, 2020, Quality, variable markups, and welfare: A quantitative general equilibrium analysis of export prices, *Journal of International Economics* 125, 103327.

Fang, Lu, Zhe Yuan, Kaifu Zhang, Dante Donati, and Miklos Sarvary, 2025, Generative AI and firm productivity: Field experiments in online retail, *arXiv*.

Felten, Edward, Manav Raj, and Robert Seamans, 2021, Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses, *Strategic Management Journal* 42, 2195–2217.

Freund, Caroline, and Diana Weinhold, 2004, The effect of the internet on international trade, *Journal of International Economics* 62, 171–189.

Goldfarb, Avi, and Daniel Trefler, 2019, Artificial intelligence and international trade, in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, ed.: *The Economics of Artificial Intelligence: An Agenda* . pp. 463–492 (University of Chicago Press).

Goldfarb, Avi, and Catherine Tucker, 2019, Digital economics, *Journal of Economic Literature* 57, 3–43.

Gozen, Suleyman Faruk, Pinar Gunes, and Mehmet Furkan Karaca, 2025, Artificial intelligence, trade, and firm dynamics, Working Paper 5016794 University of Bristol.

Hallak, Juan Carlos, and Peter K. Schott, 2011, Estimating cross-country differences in product quality, *Quarterly Journal of Economics* 126, 417–474.

Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, 2004, Export versus FDI with heterogeneous firms, *American Economic Review* 94, 300–316.

Hjort, Jonas, and Jonas Poulsen, 2019, The arrival of fast internet and employment in africa, *American Economic Review* 109, 1032–1079.

Hu, Krystal, 2023, ChatGPT sets record for fastest-growing user base: analyst note, Reuters February 2, 2023.

Hummels, David, and Peter J. Klenow, 2005, The variety and quality of a nation's exports, *American Economic Review* 95, 704–723.

Jakubik, Adam, Lorenzo Rotunno, and Alisha Saini, 2025, Foresee the unseen: Evaluating the impact of artificial intelligence on international trade, *Journal of Policy Modeling* 47, 842–861.

Johnson, Robert C., 2012, Trade and prices with heterogeneous firms, *Journal of International Economics* 86, 43–56.

Khandelwal, Amit, 2010, The long and short (of) quality ladders, *Review of Economic Studies* 77, 1450–1476.

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

Lendle, Andreas, Marcelo Olarreaga, Simon Schropp, and Pierre-Louis Vézina, 2016, There goes gravity: eBay and the death of distance, *The Economic Journal* 126, 406–441.

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

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

Noy, Shakked, and Whitney Zhang, 2023, Experimental evidence on the productivity effects of generative artificial intelligence, *Science* 381, 187–192.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, 2023, The impact of AI on developer productivity: Evidence from GitHub Copilot, *arXiv*.

Rauch, James E., 1999, Networks versus markets in international trade, *Journal of International Economics* 48, 7–35.

Schott, Peter K., 2004, Across-product versus within-product specialization in international trade, *Quarterly Journal of Economics* 119, 647–678.

Simonovska, Ina, and Michael E. Waugh, 2014, The elasticity of trade: Estimates and evidence, *Journal of International Economics* 92, 34–50.

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

Sun, Ruiqi, and Daniel Trefler, 2023, The impact of AI and cross-border data regulation on international trade in digital services: A large language model, NBER Working Paper 31925 National Bureau of Economic Research.

Verhoogen, Eric A., 2008, Trade, quality upgrading, and wage inequality in the mexican manufacturing sector, *Quarterly Journal of Economics* 123, 489–530.

A Model Equations and Equilibrium Objects

This appendix derives the main model objects used in Section 3. All integrals use the Pareto density

$$dF_j(a) = k_j \underline{a}_j^{k_j} a^{-k_j-1} da, \qquad a \geq \underline{a}_j,$$

with $k_j > \vartheta$. The restriction guarantees finite revenue integrals.

A.1 Firm-Level Revenue, Quantity, and Cutoffs

Demand for a U.S. variety in destination d is

$$q_{djt}(a,p) = D_{djt}(a)\,\xi(a)^{\sigma-1} p^{-\sigma}, \qquad \xi(a) = a^{1+\eta}.$$

With marginal cost $\tau_{dj} a/[A_j \Psi_{jt}(a)]$ and CES markup $\mu = \sigma/(\sigma-1)$, the delivered price is

$$p_{djt}(a) = \mu \tau_{dj} \frac{a}{A_j \Psi_{jt}(a)}.$$

Substituting this price into demand gives

$$q_{djt}(a) = D_{djt}(a) a^{(1+\eta)(\sigma-1)} \left(\mu \tau_{dj} \frac{a}{A_j \Psi_{jt}(a)} \right)^{-\sigma}.$$

Collecting powers of a gives

$$(1+\eta)(\sigma-1) - \sigma = \eta(\sigma-1) - 1 = \vartheta - 1,$$

so quantity can be written as

$$q_{djt}(a) = D_{djt}(a) (\mu \tau_{dj})^{-\sigma} A_j^{\sigma} \Psi_{jt}(a)^{\sigma} a^{\vartheta-1}.$$

Multiplying by $p_{djt}(a)$ gives per-firm revenue

$$r_{djt}(a) = D_{djt}(a) (\mu \tau_{dj})^{1-\sigma} A_j^{\sigma-1} \Psi_{jt}(a)^{\sigma-1} a^{\vartheta}, \qquad \vartheta \equiv \eta(\sigma-1). \tag{24}$$

For foreign destinations, $D_{djt}(a) = \Gamma_{dj,0} \exp\{-\alpha M_{dt} s_j a^{-\rho}\}$ and $\Psi_{jt}(a) = \exp\{\nu M_{Ut} s_j a^{-\rho}\}$. Hence

$$r_{djt}(a) = \Gamma_{dj,0} (\mu \tau_{dj})^{1-\sigma} A_j^{\sigma-1} a^{\vartheta} \exp\left\{ s_j [(\sigma-1)\nu M_{Ut} - \alpha M_{dt}] a^{-\rho} \right\}, \tag{25}$$

$$q_{djt}(a) = \Gamma_{dj,0} (\mu \tau_{dj})^{-\sigma} A_j^{\sigma} a^{\vartheta-1} \exp\left\{ s_j [\sigma \nu M_{Ut} - \alpha M_{dt}] a^{-\rho} \right\}. \tag{26}$$

The foreign cutoff a^*_{djt} satisfies $r_{djt}(a^*_{djt}) = \sigma f_{dj}$, which yields equation (14). The home cutoff satisfies the same zero-profit condition with $D_{Hjt} = \beta_j L_t P_{jt}^{\sigma-1}$ and no destination demand contraction.

A.2 Baseline Product-Destination Moments

At the pre-ChatGPT steady state, $M_{Ut} = M_{dt} = 0$, $\Psi_{jt}(a) = 1$, and $D_{djt}(a) = \Gamma_{dj,0}$ for foreign markets. Then

$$r_{dj,0}(a) = \Gamma_{dj,0}\,(\mu\tau_{dj})^{1-\sigma}\,A_j^{\sigma-1}\,a^{\vartheta},$$
$$q_{dj,0}(a) = \Gamma_{dj,0}\,(\mu\tau_{dj})^{-\sigma}\,A_j^{\sigma}\,a^{\vartheta-1}.$$

The revenue integral is

$$\int_{a^*_{dj,0}}^{\infty} a^{\vartheta}\,dF_j(a) = k_j \underline{a}_j^{k_j} \int_{a^*_{dj,0}}^{\infty} a^{\vartheta-k_j-1}da = \frac{k_j \underline{a}_j^{k_j} (a^*_{dj,0})^{\vartheta-k_j}}{k_j-\vartheta}.$$

The quantity integral is

$$\int_{a^*_{dj,0}}^{\infty} a^{\vartheta-1}\,dF_j(a) = k_j \underline{a}_j^{k_j} \int_{a^*_{dj,0}}^{\infty} a^{\vartheta-k_j-2}da = \frac{k_j \underline{a}_j^{k_j} (a^*_{dj,0})^{\vartheta-1-k_j}}{k_j-\vartheta+1}.$$

Multiplying by the pre-selection mass and the destination-product constants gives

$$R^x_{dj,0} = N_{j,0}\,\Gamma_{dj,0}\,(\mu\tau_{dj})^{1-\sigma}\,A_j^{\sigma-1}\,\frac{k_j \underline{a}_j^{k_j}\,(a^*_{dj,0})^{\vartheta-k_j}}{k_j-\vartheta}, \tag{27}$$

$$Q^x_{dj,0} = N_{j,0}\,\Gamma_{dj,0}\,(\mu\tau_{dj})^{-\sigma}\,A_j^{\sigma}\,\frac{k_j \underline{a}_j^{k_j}\,(a^*_{dj,0})^{\vartheta-1-k_j}}{k_j-\vartheta+1}. \tag{28}$$

The baseline export unit value is therefore

$$p^x_{dj,0} \equiv \frac{R^x_{dj,0}}{Q^x_{dj,0}} = \frac{\mu\tau_{dj}}{A_j}\,\frac{k_j-\vartheta+1}{k_j-\vartheta}\,a^*_{dj,0}. \tag{29}$$

Thus, holding the cost shifter fixed, baseline unit values are monotone in the cutoff.

A.3 Free Entry and the Home Price Index

For any market $m \in \{H\} \cup \mathcal{D}$, the operating-profit term inside free entry is

$$J_{mjt} \equiv \int_{a^*_{mjt}}^{\infty} \left[\frac{r_{mjt}(z)}{r_{mjt}(a^*_{mjt})} - 1 \right] dF_j(z).$$

At the baseline,

$$\frac{r_{mj,0}(z)}{r_{mj,0}(a^*_{mj,0})} = \left(\frac{z}{a^*_{mj,0}} \right)^{\vartheta}.$$

Hence

$$\begin{aligned} J_{mj,0} &= \int_{a^*_{mj,0}}^{\infty} \left[\left(\frac{z}{a^*_{mj,0}} \right)^{\vartheta} - 1 \right] dF_j(z) \\ &= (a^*_{mj,0})^{-\vartheta} k_j \underline{a}_j^{k_j} \int_{a^*_{mj,0}}^{\infty} z^{\vartheta - k_j - 1} dz - \left(\frac{\underline{a}_j}{a^*_{mj,0}} \right)^{k_j}. \end{aligned}$$

Using $k_j > \vartheta$,

$$J_{mj,0} = \left(\frac{\underline{a}_j}{a^*_{mj,0}} \right)^{k_j} \frac{\vartheta}{k_j - \vartheta}. \tag{30}$$

Equation (15) is then

$$\delta f_j^E = \sum_{m \in \{H\} \cup \mathcal{D}} f_{mj} J_{mjt}.$$

The home price index follows from aggregating delivered prices across surviving home varieties:

$$P_{jt}^{1-\sigma} = N_{jt}\, \mu^{1-\sigma} A_j^{\sigma-1} k_j \underline{a}_j^{k_j} \int_{a_{jt}^{H*}}^{\infty} a^{\vartheta - k_j - 1} \exp\{(\sigma - 1)\nu M_{Ut} s_j a^{-\rho}\} da. \tag{31}$$

Once the foreign cutoffs are known, the free-entry condition determines the home cutoff. The home zero-profit condition then returns the home price index, and (31) returns the pre-selection mass.

B Proofs of the Model Propositions

All derivatives in this section are evaluated at the pre-ChatGPT steady state. At that point $M_{Ut} = M_{dt} = 0$, $\Psi_{jt}(a) = 1$, and the foreign demand shifter equals $\Gamma_{dj,0}$.

B.1 A Pareto Differentiation Identity

The proofs repeatedly use the same derivative. For $\chi_j < k_j$, define

$$\Phi_\chi(a, X) \equiv \int_a^\infty z^{\chi_j - k_j - 1} \exp\{s_j X z^{-\rho}\} dz.$$

At $X = 0$,

$$\Phi_\chi(a, 0) = \frac{a^{\chi_j - k_j}}{k_j - \chi_j}.$$

Leibniz's rule gives the cutoff component

$$\left.\frac{\partial \log \Phi_\chi}{\partial \log a}\right|_{X=0} = \frac{-a^{\chi_j - k_j}}{a^{\chi_j - k_j}/(k_j - \chi_j)} = \chi_j - k_j.$$

The shock component is

$$\left.\frac{\partial \log \Phi_\chi}{\partial X}\right|_{X=0} = s_j \frac{\int_a^\infty z^{\chi_j - k_j - 1 - \rho} dz}{\int_a^\infty z^{\chi_j - k_j - 1} dz} = s_j a^{-\rho} \frac{k_j - \chi_j}{k_j + \rho - \chi_j}.$$

Therefore

$$d \log \Phi_\chi = (\chi_j - k_j) d \log a + s_j a^{-\rho} \frac{k_j - \chi_j}{k_j + \rho - \chi_j} dX. \tag{32}$$

B.2 Proof of Proposition 1

Foreign cutoff. Taking logs of equation (14) gives

$$\log\left[\frac{\Gamma_{dj,0}(\mu \tau_{dj})^{1-\sigma} A_j^{\sigma-1}}{\sigma}\right] + \vartheta \log a^*_{djt} + s_j (a^*_{djt})^{-\rho}[(\sigma - 1)\nu M_{Ut} - \alpha M_{dt}] = \log f_{dj}.$$

The bracketed shock term is zero at the baseline, so the derivative of $(a^*_{djt})^{-\rho}$ does not enter at first order. Differentiating gives

$$\vartheta \, d \log a^*_{djt} + s_j (a^*_{dj,0})^{-\rho} \left[(\sigma - 1)\nu \, dM_{Ut} - \alpha \, dM_{dt}\right] = 0.$$

Solving delivers equation (17).

Export unit value. Write aggregate foreign revenue and quantity as

$$R^x_{djt} = N_{jt}\Gamma_{dj,0}(\mu\tau_{dj})^{1-\sigma}A_j^{\sigma-1}k_j\underline{a}_j^{k_j}\Phi^R_{djt},$$
$$Q^x_{djt} = N_{jt}\Gamma_{dj,0}(\mu\tau_{dj})^{-\sigma}A_j^{\sigma}k_j\underline{a}_j^{k_j}\Phi^Q_{djt},$$

where

$$\Phi^R_{djt} = \int_{a^*_{djt}}^{\infty} a^{\vartheta-k_j-1}\exp\big\{s_j[(\sigma-1)\nu M_{Ut} - \alpha M_{dt}]a^{-\rho}\big\}\,da,$$
$$\Phi^Q_{djt} = \int_{a^*_{djt}}^{\infty} a^{\vartheta-k_j-2}\exp\big\{s_j[\sigma\nu M_{Ut} - \alpha M_{dt}]a^{-\rho}\big\}\,da.$$

The mass N_{jt} cancels in $\log p^x_{djt} = \log R^x_{djt} - \log Q^x_{djt}$. Leibniz differentiation at the baseline gives

$$d\log\Phi^R_{djt} = (\vartheta-k_j)\,d\log a^*_{djt} + \frac{(k_j-\vartheta)\,s_j(a^*_{dj,0})^{-\rho}}{k_j+\rho-\vartheta}\big[(\sigma-1)\nu\,dM_{Ut} - \alpha\,dM_{dt}\big],$$
$$d\log\Phi^Q_{djt} = (\vartheta-k_j-1)\,d\log a^*_{djt} + \frac{(k_j-\vartheta+1)\,s_j(a^*_{dj,0})^{-\rho}}{k_j+\rho-\vartheta+1}\big[\sigma\nu\,dM_{Ut} - \alpha\,dM_{dt}\big].$$

The first line applies (32) with $\chi_j = \vartheta$ and $dX = (\sigma-1)\nu\,dM_{Ut} - \alpha\,dM_{dt}$. The second applies it with $\chi_j = \vartheta - 1$ and $dX = \sigma\nu\,dM_{Ut} - \alpha\,dM_{dt}$. Subtracting the two lines gives

$$\begin{aligned} d\log p^x_{djt} = {} & d\log a^*_{djt} \\ & + s_j(a^*_{dj,0})^{-\rho}\left[\frac{k_j-\vartheta}{k_j+\rho-\vartheta}\big((\sigma-1)\nu\,dM_{Ut} - \alpha\,dM_{dt}\big)\right. \\ & \qquad \left. -\frac{k_j-\vartheta+1}{k_j+\rho-\vartheta+1}\big(\sigma\nu\,dM_{Ut} - \alpha\,dM_{dt}\big)\right]. \end{aligned}$$

Substituting (17) yields the coefficient on destination adoption:

$$\alpha\left[\frac{1}{\vartheta} - \frac{k_j-\vartheta}{k_j+\rho-\vartheta} + \frac{k_j-\vartheta+1}{k_j+\rho-\vartheta+1}\right] = \alpha\left[\frac{1}{\vartheta} + \frac{\rho}{(k_j+\rho-\vartheta)(k_j+\rho-\vartheta+1)}\right],$$

where the equality follows by setting $H_j \equiv k_j+\rho-\vartheta$ and using

$$-\frac{H_j-\rho}{H_j} + \frac{H_j+1-\rho}{H_j+1} = \frac{\rho}{H_j(H_j+1)}.$$

The coefficient on U.S. adoption enters with the opposite sign. Thus (18) holds, with

$$K_j^d = \alpha\left[\frac{1}{\vartheta} + \frac{\rho}{(k_j+\rho-\vartheta)(k_j+\rho-\vartheta+1)}\right] > 0, \tag{33}$$

$$K_j^U = \nu\left[\frac{\sigma-1}{\vartheta} - \frac{(k_j-\vartheta)(\sigma-1)}{k_j+\rho-\vartheta} + \frac{(k_j-\vartheta+1)\sigma}{k_j+\rho-\vartheta+1}\right] > 0. \tag{34}$$

Export quantity. Holding N_{jt} fixed, the quantity response comes from Φ_{djt}^Q. Let $\chi_j \equiv \vartheta - 1$. Then

$$d\log\Phi_{djt}^Q = -(k_j-\chi_j)\,d\log a_{djt}^* + \frac{(k_j-\chi_j)\,s_j(a_{dj,0}^*)^{-\rho}}{k_j+\rho-\chi_j}\big[\sigma\nu\,dM_{Ut} - \alpha\,dM_{dt}\big].$$

Substituting (17) gives

$$d\log Q_{djt}^x\,\big|_{dN=0} = s_j(a_{dj,0}^*)^{-\rho}\Bigg\{-\alpha(k_j-\chi_j)\left[\frac{1}{\vartheta}+\frac{1}{k_j+\rho-\chi_j}\right]dM_{dt} + \nu(k_j-\chi_j)\left[\frac{\sigma-1}{\vartheta}+\frac{\sigma}{k_j+\rho-\chi_j}\right]dM_{Ut}\Bigg\}.$$

This is (19), with

$$K_j^{Q,d} = \alpha\,\frac{(k_j-\chi_j)(\vartheta+k_j+\rho-\chi_j)}{\vartheta\,(k_j+\rho-\chi_j)} > 0, \tag{35}$$

$$K_j^{Q,U} = \nu\left[\frac{(k_j-\chi_j)(\sigma-1)}{\vartheta} + \frac{(k_j-\chi_j)\sigma}{k_j+\rho-\chi_j}\right] > 0. \tag{36}$$

B.3 Proof of Proposition 2

For compactness define

$$d\mathcal{X}_{djt} \equiv (\sigma-1)\nu\,dM_{Ut} - \alpha\,dM_{dt}, \qquad d\mathcal{X}_{Hjt} \equiv (\sigma-1)\nu\,dM_{Ut}.$$

For any market m, the free-entry integral can be written as

$$J_m(a_{mjt}^*, \mathcal{X}_{mjt}) = \int_{a_{mjt}^*}^{\infty}\left[\left(\frac{z}{a_{mjt}^*}\right)^{\vartheta}\exp\{s_j\mathcal{X}_{mjt}(z^{-\rho}-(a_{mjt}^*)^{-\rho})\}-1\right]dF_j(z).$$

At $\mathcal{X}_{mjt} = 0$, this integral equals (30). Its cutoff derivative is immediate because $J_m(a^*, 0)$ is proportional to $(a^*)^{-k_j}$:

$$d \log J_m \big|_{\mathcal{X}=0,\, d\mathcal{X}=0} = -k_j \, d \log a^*_{mjt}.$$

The derivative with respect to $\mathcal{X}_{mjt}$ is

$$\begin{aligned}
\left.\frac{\partial J_m}{\partial \mathcal{X}}\right|_{\mathcal{X}=0} &= s_j \int_{a^*}^{\infty} \left(\frac{z}{a^*}\right)^{\vartheta} \left[z^{-\rho} - (a^*)^{-\rho}\right] dF_j(z) \\
&= s_j k_j \underline{a}_j^{k_j} (a^*)^{-\vartheta} \left[\int_{a^*}^{\infty} z^{\vartheta-\rho-k_j-1} dz - (a^*)^{-\rho} \int_{a^*}^{\infty} z^{\vartheta-k_j-1} dz\right] \\
&= -s_j k_j \underline{a}_j^{k_j} (a^*)^{-k_j-\rho} \frac{\rho}{(k_j + \rho - \vartheta)(k_j - \vartheta)}.
\end{aligned}$$

Dividing by $J_m(a^*, 0) = (\underline{a}_j / a^*)^{k_j} \vartheta / (k_j - \vartheta)$ gives

$$\left.\frac{1}{J_m} \frac{\partial J_m}{\partial \mathcal{X}}\right|_{\mathcal{X}=0} = -s_j (a^*)^{-\rho} \frac{k_j \rho}{\vartheta(k_j + \rho - \vartheta)}.$$

Thus the first-order expansion is

$$d \log J_m = -k_j \, d \log a^*_{mjt} - C_j^E \, s_j (a^*_{mj,0})^{-\rho} d\mathcal{X}_{mjt}, \qquad C_j^E \equiv \frac{k_j \rho}{\vartheta(k_j + \rho - \vartheta)}. \tag{37}$$

Let $R^H_{j,0}$ denote baseline home-market revenue, let $R^X_{j,0} \equiv \sum_{d \in \mathcal{D}} R^x_{dj,0}$, and define

$$\Omega^H_{j,0} \equiv \frac{R^H_{j,0}}{R^H_{j,0} + R^X_{j,0}}, \qquad \Omega^x_{dj,0} \equiv \frac{R^x_{dj,0}}{R^H_{j,0} + R^X_{j,0}}.$$

These are the free-entry weights. To see this, use the zero-profit condition $r_{mj,0}(a^*_{mj,0}) = \sigma f_{mj}$ and the baseline revenue integral. For market m,

$$f_{mj} J_{mj,0} = \frac{r_{mj,0}(a^*_{mj,0})}{\sigma} \left(\frac{\underline{a}_j}{a^*_{mj,0}}\right)^{k_j} \frac{\vartheta}{k_j - \vartheta} = \frac{\vartheta}{\sigma k_j} \frac{R^m_{j,0}}{N_{j,0}}.$$

The factor $\vartheta / (\sigma k_j N_{j,0})$ is common across markets, so differentiating free entry is equivalent to

$$\Omega^H_{j,0} d \log J_H + \sum_{d \in \mathcal{D}} \Omega^x_{dj,0} d \log J_d = 0.$$

For a foreign destination, (17) implies

$$d\log a^*_{djt} = -\frac{s_j(a^*_{dj,0})^{-\rho}}{\vartheta} d\mathcal{X}_{djt}.$$

Substituting this into (37) gives the foreign contribution

$$d\log J_d = s_j(a^*_{dj,0})^{-\rho}\left(\frac{k_j}{\vartheta} - C^E_j\right) d\mathcal{X}_{djt} = C^X_j s_j(a^*_{dj,0})^{-\rho} d\mathcal{X}_{djt},$$

where

$$C^X_j \equiv \frac{k_j(k_j-\vartheta)}{\vartheta(k_j+\rho-\vartheta)}.$$

The home contribution is

$$d\log J_H = -k_j d\log a^{H*}_{jt} - C^E_j s_j(a^*_{Hj,0})^{-\rho} d\mathcal{X}_{Hjt}.$$

Combining the home and foreign contributions gives

$$k_j\Omega^H_{j,0}\, d\log a^{H*}_{jt} = s_j\left[C^X_j \sum_{d\in\mathcal{D}} \Omega^x_{dj,0}(a^*_{dj,0})^{-\rho} d\mathcal{X}_{djt} - C^E_j \Omega^H_{j,0}(a^*_{Hj,0})^{-\rho} d\mathcal{X}_{Hjt}\right], \tag{38}$$

The home zero-profit condition is

$$\log(\beta_j L_t) + (\sigma-1)\log P_{jt} + \log(\mu^{1-\sigma}A_j^{\sigma-1}/\sigma) + \vartheta\log a^{H*}_{jt} + s_j(a^{H*}_{jt})^{-\rho}\mathcal{X}_{Hjt} = \log f^H_j.$$

At the baseline, the derivative of $(a^{H*}_{jt})^{-\rho}$ again drops out because $\mathcal{X}_{Hjt} = 0$. Thus

$$(\sigma-1)d\log P_{jt} = -\vartheta\, d\log a^{H*}_{jt} - s_j(a^*_{Hj,0})^{-\rho} d\mathcal{X}_{Hjt}, \tag{39}$$

The home price-index equation can be written as

$$(1-\sigma)d\log P_{jt} = d\log N_{jt} + d\log\Phi^H_{jt},$$

where

$$\Phi^H_{jt} = \int_{a^{H*}_{jt}}^{\infty} a^{\vartheta-k_j-1}\exp\{s_j\mathcal{X}_{Hjt}a^{-\rho}\}da.$$

Using (32) with $\chi_j = \vartheta$ gives

$$d\log\Phi^H_{jt} = (\vartheta-k_j)d\log a^{H*}_{jt} + \frac{k_j-\vartheta}{k_j+\rho-\vartheta} s_j(a^*_{Hj,0})^{-\rho} d\mathcal{X}_{Hjt}.$$

Combining this expression with (39) yields

$$d\log N_{jt} = k_j\, d\log a_{jt}^{H*} + \frac{\rho}{k_j+\rho-\vartheta} s_j (a_{Hj,0}^*)^{-\rho} d\mathcal{X}_{Hjt}. \tag{40}$$

Combining (38) and (40) gives

$$d\log N_{jt} = s_j \left[\frac{C_j^X}{\Omega_{j,0}^H} \sum_{d\in\mathcal{D}} \Omega_{dj,0}^x (a_{dj,0}^*)^{-\rho} d\mathcal{X}_{djt} - \Theta_j^H (a_{Hj,0}^*)^{-\rho} d\mathcal{X}_{Hjt} \right],$$

where

$$\Theta_j^H = \frac{\rho(k_j-\vartheta)}{\vartheta(k_j+\rho-\vartheta)} > 0.$$

Using $\Omega_{dj,0}^x/\Omega_{j,0}^H = \omega_{dj,0}^N R_{j,0}^X/R_{j,0}^H$ and expanding $d\mathcal{X}_{djt}$ yields equation (20), with

$$\begin{aligned}
\Lambda_j^N &= \frac{k_j(k_j-\vartheta)}{\vartheta(k_j+\rho-\vartheta)} \frac{R_{j,0}^X}{R_{j,0}^H} > 0, \\
B_j^U &= (\sigma-1)\nu \left[\sum_{d\in\mathcal{D}} \omega_{dj,0}^N (a_{dj,0}^*)^{-\rho} \Lambda_j^N - \Theta_j^H (a_{Hj,0}^*)^{-\rho} \right], \\
B_{dj}^D &= \alpha\, \omega_{dj,0}^N \Lambda_j^N (a_{dj,0}^*)^{-\rho} > 0.
\end{aligned}$$

The strict positivity of Λ_j^N, Θ_j^H, and B_{dj}^D follows from $k_j > \vartheta$, $\rho > 0$, $\alpha > 0$, and positive baseline home and foreign revenues.

C Solving the Equilibrium

This section gives the equilibrium-solving algebra used in the quantitative analysis. The main text denotes the pre-selection mass by N_{jt}. In the quantitative implementation this same object is stored as the draw mass M_{jt}; the notation below follows the main text.

C.1 Integral Blocks

Define the foreign and home shock indexes

$$\mathcal{Z}_{djt} \equiv s_j[(\sigma-1)\nu M_{Ut} - \alpha M_{dt}], \qquad \mathcal{Z}^H_{jt} \equiv s_j(\sigma-1)\nu M_{Ut}.$$

For any cutoff a and any shock index $\mathcal{Z}$, let $B(a) = \mathcal{Z}a^{-\rho}$. The Pareto integrals are evaluated with

$$\int_a^\infty z^{p-1}\exp\{\mathcal{Z}z^{-\rho}\}dz = \frac{a^p}{\rho}\mathcal{S}(B(a);\gamma_p), \qquad \gamma_p = -\frac{p}{\rho} > 0, \tag{41}$$

where

$$\mathcal{S}(B;\gamma) \equiv \sum_{n=0}^{\infty}\frac{B^n}{(\gamma+n)n!}. \tag{42}$$

The identity follows from the change of variables $u = (a/z)^\rho$. Then $z = au^{-1/\rho}$, $dz = -(a/\rho)u^{-1/\rho-1}du$, and

$$z^{p-1}dz = -\frac{a^p}{\rho}u^{-p/\rho-1}du = -\frac{a^p}{\rho}u^{\gamma_p-1}du.$$

The limits $z = a$ and $z = \infty$ become $u = 1$ and $u = 0$. Hence

$$\int_a^\infty z^{p-1}\exp\{\mathcal{Z}z^{-\rho}\}dz = \frac{a^p}{\rho}\int_0^1 u^{\gamma_p-1}\exp\{B(a)u\}du.$$

Expanding the exponential gives (42). The series is absolutely convergent for finite B.

For the free-entry block define

$$\gamma_j \equiv \frac{k_j - \vartheta}{\rho}.$$

The market-m free-entry integral, with $m = d$ or $m = H$, is

$$J_{mjt}(a^*) = \left(\frac{a_j}{a^*}\right)^{k_j}\left[\frac{k_j\exp\{-B_m(a^*)\}}{\rho}\mathcal{S}(B_m(a^*);\gamma_j) - 1\right], \tag{43}$$

where $B_d(a^*) = \mathcal{Z}_{djt}(a^*)^{-\rho}$ and $B_H(a^*) = \mathcal{Z}^H_{jt}(a^*)^{-\rho}$. At $M_{Ut} = M_{dt} = 0$, equation (43) collapses to (30).

The home price-index integral is

$$\mathcal{I}^{P}_{Hjt}(a^{H*}) = \frac{(a^{H*})^{\vartheta-k_j}}{\rho}\mathcal{S}(B_H(a^{H*});\gamma_j). \tag{44}$$

Foreign revenue and quantity integrals are

$$\mathcal{I}^{R}_{djt}(a^*) = \frac{(a^*)^{\vartheta-k_j}}{\rho}\mathcal{S}\left(\mathcal{Z}_{djt}(a^*)^{-\rho};\frac{k_j-\vartheta}{\rho}\right), \tag{45}$$

$$\mathcal{I}^{Q}_{djt}(a^*) = \frac{(a^*)^{\vartheta-k_j-1}}{\rho}\mathcal{S}\left(s_j[\sigma\nu M_{Ut}-\alpha M_{dt}](a^*)^{-\rho};\frac{k_j-\vartheta+1}{\rho}\right). \tag{46}$$

C.2 Solving Algebra

Given calibrated primitives

$$\{A_j,\Gamma_{dj,0},\tau_{dj},f_{dj},f^H_j,f^E_j,\beta_j,\underline{a}_j,s_j\}_{d,j}$$

and a shock path $\{M_{Ut},M_{dt}\}_d$, the equilibrium is computed as follows.

1. *Foreign cutoffs.* For each product-destination market, solve

$$\Gamma_{dj,0}(\mu\tau_{dj})^{1-\sigma}A_j^{\sigma-1}(a^*_{djt})^{\vartheta}\exp\{\mathcal{Z}_{djt}(a^*_{djt})^{-\rho}\} = \sigma f_{dj}. \tag{47}$$

   This is one scalar equation for each product and destination. If the left side at $\underline{a}_j$ exceeds σf_{dj}, set $a^*_{djt}=\underline{a}_j$. Otherwise solve by bisection on $[\underline{a}_j,\bar{a}]$, expanding $\bar{a}$ until the root is bracketed.

2. *Home cutoff.* Given the foreign cutoffs, solve the free-entry condition

$$\delta f^E_j = f^H_j J_{Hjt}(a^{H*}_{jt}) + \sum_{d\in\mathcal{D}} f_{dj}J_{djt}(a^*_{djt}) \tag{48}$$

   for a^{H*}_{jt}. This is one scalar equation per product. If the right side at $\underline{a}_j$ is below δf^E_j, the product is inactive under that shock path.

3. *Home price index.* For active products, compute P_{jt} from the home zero-profit condition:

$$P^{\sigma-1}_{jt} = \frac{\sigma f^H_j}{\beta_j L_t\,\mu^{1-\sigma}A_j^{\sigma-1}(a^{H*}_{jt})^{\vartheta}\exp\{\mathcal{Z}^H_{jt}(a^{H*}_{jt})^{-\rho}\}}. \tag{49}$$

4. *Pre-selection mass.* Compute N_{jt} from the home price index:

$$N_{jt} = \frac{P_{jt}^{1-\sigma}}{\mu^{1-\sigma} A_j^{\sigma-1} k_j \underline{a}_j^{k_j} \mathcal{I}_{Hjt}^{P}(a_{jt}^{H*})}. \tag{50}$$

5. *Export observables.* For every active export market,

$$R_{djt}^{x} = N_{jt} \Gamma_{dj,0} (\mu \tau_{dj})^{1-\sigma} A_j^{\sigma-1} k_j \underline{a}_j^{k_j} \mathcal{I}_{djt}^{R}(a_{djt}^{*}), \tag{51}$$

$$Q_{djt}^{x} = N_{jt} \Gamma_{dj,0} (\mu \tau_{dj})^{-\sigma} A_j^{\sigma} k_j \underline{a}_j^{k_j} \mathcal{I}_{djt}^{Q}(a_{djt}^{*}), \tag{52}$$

$$p_{djt}^{x} = R_{djt}^{x} / Q_{djt}^{x}.$$

6. *Labor-market clearing.* Homogeneous-good production absorbs the labor not used in differentiated production, market-access fixed costs, and entry. The labor-market constraint therefore pins homogeneous-good output and closes the equilibrium.

The algebra is sequential because the foreign cutoff equations are independent across destinations conditional on primitives and shocks; the home free-entry equation then pins the home cutoff; and the home zero-profit and price-index equations return P_{jt} and N_{jt} in closed form.

D Calibration Details

This appendix gives the calibration algebra used in Section 7.1. The procedure has four steps. First, we derive the pre-AI closed-form equilibrium that expresses every endogenous baseline object as an explicit function of the time-invariant fundamentals and the trade cost. Second, destination iceberg costs are recovered from pre-period gravity variation, using the gravity reduced form implied by that closed form. Third, for each trial value of η, the time-invariant fundamentals are backed out from observed revenues, unit values, establishment counts, survival rates, and exporter shares using the same closed-form relations. Fourth, the AI-shock parameters are estimated by matching simulated reduced-form moments to their empirical counterparts.

D.1 Pre-AI Closed-Form Equilibrium

The baseline period sets the AI shocks to zero, so the shock factor $\exp(\Lambda_{dj}(a^*)^{-\rho})$ in (25) collapses to one and every baseline integral admits a closed form. We use this subsection to collect the two pre-AI identities that the rest of the appendix calls on repeatedly: the closed-form foreign cutoff $a^*_{dj,0}$ as a function of the time-invariant fundamentals, and the resulting bilateral revenue $R^x_{dj,0}$ as a power of τ_{dj}.

Foreign cutoff. The foreign zero-profit condition is $r_{dj,0}(a^*_{dj,0}) = \sigma f_{dj}$, with the firm-level revenue $r_{dj,0}(a) = \Gamma_{dj,0}(\mu\tau_{dj})^{1-\sigma} A_j^{\sigma-1} a^{\vartheta}$ from the model section. Solving for the cutoff,

$$a^*_{dj,0} = \left(\frac{\sigma f_{dj}}{\Gamma_{dj,0}}\right)^{1/\vartheta} (\mu\tau_{dj})^{(\sigma-1)/\vartheta} A_j^{-(\sigma-1)/\vartheta}. \tag{53}$$

Equation (53) is the unique solution under the maintained convergence condition $k_j > \vartheta$ from Section 3; it is monotone in τ_{dj} with elasticity $(\sigma-1)/\vartheta$.

Bilateral revenue as a power of τ_{dj}. Aggregating firm-level revenue over the Pareto draw distribution above $a^*_{dj,0}$ delivers the baseline bilateral revenue (27),

$$R^x_{dj,0} = N_{j,0}\,\Gamma_{dj,0}\,(\mu\tau_{dj})^{1-\sigma}\,A_j^{\sigma-1}\,\frac{k_j \underline{a}_j^{k_j}\,(a^*_{dj,0})^{\vartheta-k_j}}{k_j-\vartheta}.$$

Substituting the closed-form cutoff (53) into this expression eliminates the endogenous $a^*_{dj,0}$

on the right-hand side and groups all τ_{dj} exponents:

$$R^x_{dj,0} = N_{j,0}\,\frac{k_j a_j^{k_j}}{k_j-\vartheta}\,(\mu\tau_{dj})^{1-\sigma}\,A_j^{\sigma-1}\,\Gamma_{dj,0}\,\left[\left(\frac{\sigma f_{dj}}{\Gamma_{dj,0}}\right)^{1/\vartheta}(\mu\tau_{dj})^{(\sigma-1)/\vartheta}A_j^{-(\sigma-1)/\vartheta}\right]^{\vartheta-k_j}$$
$$= C_{dj}\,(\mu\tau_{dj})^{-(\sigma-1)k_j/\vartheta}, \tag{54}$$

where the constant C_{dj} collects the τ-independent terms,

$$C_{dj} \equiv \frac{N_{j,0}\,k_j a_j^{k_j}}{k_j-\vartheta}\,A_j^{(\sigma-1)k_j/\vartheta}\,\Gamma_{dj,0}^{k_j/\vartheta}\,(\sigma f_{dj})^{1-k_j/\vartheta}.$$

Using $\vartheta = \eta(\sigma-1)$ from the preset constants, $(\sigma-1)k_j/\vartheta = k_j/\eta$, so the iceberg elasticity of bilateral revenue is exactly $-k_j/\eta$:

$$R^x_{dj,0} \propto \tau_{dj}^{-k_j/\eta}. \tag{55}$$

This is the closed-form identity that delivers the trade-cost elasticity used in Section 7.1 for both the destination iceberg gravity regression of Section D.2 and the back-substitution that anchors the closed-form fundamentals of Section D.3.

D.2 Destination Iceberg Costs

Summing (54) across products in destination d and taking logs delivers the gravity reduced form

$$\log X_{d,0} = \alpha_0 \;-\; \tfrac{k_j}{\eta}\log\tau_d \;+\; \Phi(Y_d,\mathbf{X}_d) \;+\; u_d, \qquad X_{d,0} \equiv \textstyle\sum_j R^x_{dj,0}, \tag{56}$$

where $\Phi(Y_d,\mathbf{X}_d)$ absorbs destination size and the product aggregator of the C_{dj} that is uncorrelated with τ_d under the maintained separability of $\Gamma_{dj,0}$ into a destination size component and a product preference component. The proxy for $\log\tau_d$ is a CEPII Gravity vector of bilateral resistance variables; the regression is

$$\log X_{d,0} = \gamma_0 + \gamma_1\log\mathrm{dist}_d + \gamma_2\mathrm{lang}_d + \gamma_3\mathrm{colony}_d + \gamma_4\mathrm{fta}_d + Z_d'\gamma_Z + u_d, \tag{57}$$

with Z_d containing log GDP per capita, log population, WTO membership, and EU membership. The fitted bilateral-resistance component is

$$\hat{g}_d \equiv \hat{\gamma}_1\log\mathrm{dist}_d + \hat{\gamma}_2\mathrm{lang}_d + \hat{\gamma}_3\mathrm{colony}_d + \hat{\gamma}_4\mathrm{fta}_d,$$

and the implied iceberg follows from rearranging the model coefficient $-k_j/\eta$ in (56):

$$\log \tilde{\tau}_d = -\frac{\hat{g}_d}{k_j/\eta}.$$

The level normalization sets Canada as the reference destination,

$$\hat{\tau}_d = \frac{\tilde{\tau}_d}{\tilde{\tau}_{\text{CAN}}}, \tag{58}$$

so that $\hat{\tau}_{\text{CAN}} = 1$.

The industry component $\hat{\tau}_j$ is estimated from the normalized revenue–labor ratio in the ASM. Let $S_{j,0} = R^H_{j,0} + \sum_d R^x_{dj,0}$ denote total sales of industry j, and define the home and export revenue shares

$$\phi^H_j = \frac{R^H_{j,0}}{S_{j,0}}, \qquad \phi^x_{dj} = \frac{R^x_{dj,0}}{S_{j,0}}.$$

Given the destination icebergs from the gravity step, the destination component of the industry-average iceberg is the revenue-weighted harmonic mean

$$\bar{\tau}^d_j = \left(\phi^H_j + \sum_d \frac{\phi^x_{dj}}{\hat{\tau}_d} \right)^{-1}. \tag{59}$$

Let r_j be the industry's share of aggregate manufacturing revenue and ℓ_j its share of aggregate manufacturing labor cost. The model-implied normalized revenue–labor ratio satisfies

$$\frac{r_j}{\ell_j} = \frac{\sigma \mu \bar{\tau}_j}{\sigma + \mu \bar{\tau}_j}, \qquad \bar{\tau}_j = \tau_j \bar{\tau}^d_j. \tag{60}$$

Solving (60) gives

$$\hat{\bar{\tau}}_j = \frac{\sigma (r_j/\ell_j)}{\mu \{\sigma - (r_j/\ell_j)\}}, \tag{61}$$

defined for $0 < r_j/\ell_j < \sigma$. The industry multiplier and the bilateral iceberg are then

$$\hat{\tau}_j = \frac{\hat{\bar{\tau}}_j}{\bar{\tau}^d_j}, \qquad \hat{\tau}_{dj} = \hat{\tau}_d \, \hat{\tau}_j. \tag{62}$$

D.3 Calibration of Time-Invariant Fundamentals

This subsection has two parts. We first describe how each pre-period data input is constructed and aligned to the HS10 product unit used by the model; we then state the closed-form mapping that, given a candidate η, recovers the time-invariant fundamentals

$\{A_j, \Gamma_{dj,0}, f_{dj}, f_j^H, \beta_j, f_j^E\}$ from those inputs.

U.S. export observations. The raw inputs are the U.S. Census Bureau Schedule B HS10–country–month export records. The pre-period window is the $T_0 = 47$ months from 2019m1 through 2022m11. For each observed HS10 product j and foreign destination $d \in \mathcal{D}$, the pre-period objects entering the calibration are the average monthly revenue, the average monthly quantity, and the implied unit value,

$$R^x_{dj,0} = \frac{1}{T_0}\sum_{t=1}^{T_0} \text{value}_{djt}, \qquad Q^x_{dj,0} = \frac{1}{T_0}\sum_{t=1}^{T_0} \text{quantity}_{djt}, \qquad p^x_{dj,0} = \frac{\sum_{t=1}^{T_0} \text{value}_{djt}}{\sum_{t=1}^{T_0} \text{quantity}_{djt}}. \quad (63)$$

Destination-level pre-period exports $X_{d,0} = \sum_j R^x_{dj,0}$ enter the $\hat{\tau}_d$ gravity step in Section D.2.

Home revenue. Home revenue is U.S.-producer domestic sales. It is constructed from the Census Bureau Annual Survey of Manufactures shipments averaged across the three pre-period years 2019–2021, converted from annual to monthly frequency, and netted of monthly pre-period exports:

$$R^H_{n,0} = \frac{\overline{\text{ASM shipments}}_{n,\,2019:2021}}{12} - \sum_{j \in G(n)} \sum_{d \in \mathcal{D}} R^x_{dj,0}, \quad (64)$$

where $G(n)$ is the set of HS10 products mapped to NAICS industry n by the Pierce–Schott NAICS–HS10 concordance. The annual ASM shipments figure is in dollars and is reported at the industry level; the division by twelve and the subtraction of monthly pre-period exports ensure that $R^H_{n,0}$ has the same unit and frequency as the product-destination revenue inputs entering the model.

Home establishment count. Home establishment counts $N^H_{n,0}$ are taken from the Census Business Dynamics Statistics establishment files at the NAICS industry level for the pre-period window. In the model, the pre-selection mass $N_{j,0} = N^H_{j,0}/s^H_j$ is the mass of *potential* draws and is not separately observed; the observed count $N^H_{j,0}$ in the data is interpreted as the mass of *successful* home producers, which is what enters (72).

Successful-entry probability. The model parameter s^H_j is the probability that a potential entrant becomes an employer producer in industry j. We measure it at the NAICS industry level from the Census Business Formation Statistics. The baseline empirical coun-

terpart is the four-quarter application-to-formation rate,

$$s_n^H = \frac{\text{BF4Q}_n}{\text{HBA}_n}, \tag{65}$$

where HBA_n is the pre-period stock of high-propensity business applications and BF4Q_n is the count of employer business formations within four quarters of an application. We use $\text{BF8Q}_n/\text{HBA}_n$ (an eight-quarter formation window) and $\text{BF4Q}_n/\text{WBA}_n$ (applications with planned wages as the entrant pool) as robustness measures. Industries with $s_n^H \leq 0$ or $s_n^H > 1$ are excluded from the calibration and reported; the inversion requires $0 < s_j^H \leq 1$ to deliver a finite home cutoff in (69).

Exporter share. The exporter share π_j is the conditional probability that an active producer in industry j exports to at least one destination. It is measured at the NAICS-3 level from the 2022 Census exporter survey as

$$\pi_n = \frac{\text{exporting firms in industry } n}{\text{all firms in industry } n}. \tag{66}$$

Because π_j is a conditional probability among home producers, it identifies only the relative cutoff of the easiest foreign cutoff $\hat{a}^*_{\text{easy},j,0}$ relative to the home cutoff $\hat{a}^{H*}_{j,0}$; combined with s_j^H, the joint pair places both cutoffs on the absolute Pareto support $[\underline{a}_j, \infty)$.

Aggregate income. The aggregate income normalizer entering $\beta_j = R^H_{j,0}/L_0$ is the pre-period U.S. Bureau of Economic Analysis monthly personal-income series `PI`, averaged across the pre-period window. Using L_0 rather than industry-specific labor input keeps β_j comparable across industries.

NAICS-to-HS10 allocation and sample selection. Home revenue $R^H_{n,0}$, the establishment count $N^H_{n,0}$, the successful-entry probability s_n^H, and the exporter share π_n are observed at the NAICS industry level rather than at HS10. Each industry-level input $X_n \in \{R^H_{n,0}, N^H_{n,0}, s_n^H, \pi_n\}$ is allocated to the HS10 products in $G(n)$ by an equal-weighted split,

$$X_j = \frac{X_n}{|G(n)|}, \qquad j \in G(n), \tag{67}$$

which preserves the industry total $\sum_{j \in G(n)} X_j = X_n$ and is invariant to the choice of HS10 grouping within an industry. The HS10 sample retained for the inversion is the intersection of the Schedule B export panel with the NAICS industries for which $R^H_{n,0} > 0$, $N^H_{n,0} > 0$, and $0 < s_n^H, \pi_n \leq 1$. HS10 products that map to industries failing any of these requirements, or

that have no foreign export records over the pre-period window, are excluded and reported.

Closed-form solution given η. Fix a candidate value of η. The trade elasticity restriction $k_j/\eta = 4.25$ from Section 7.1 pins the Pareto shape and the composite revenue exponent:

$$k_j = 4.25\,\eta, \qquad \vartheta = \eta(\sigma - 1).$$

The remaining steps walk through the calibration one parameter at a time, in the order they are needed by later expressions. Each step identifies the model object on the left-hand side from one zero-profit, free-entry, or Pareto-survival condition stated elsewhere in the model section; the right-hand side is then expressed in terms of observed data and quantities recovered in previous steps.

Step 1: unscaled cutoff $\tilde{a}^*_{dj,0}$ from the unit value. The model's baseline unit-value identity (29) states $p^x_{dj,0} = (\mu\tau_{dj}/A_j)\,T_j\,a^*_{dj,0}$ with $T_j \equiv (k_j - \vartheta + 1)/(k_j - \vartheta)$ the Pareto-truncation multiplier obtained when the master integral (41) is evaluated at exponent $p = -1$. Solving for the cutoff and absorbing the unknown A_j into the auxiliary object $\tilde{a}^*_{dj,0} \equiv a^*_{dj,0}/A_j$ delivers

$$\tilde{a}^*_{dj,0} = \frac{p^x_{dj,0}}{\mu\hat{\tau}_{dj}T_j}, \tag{68}$$

which depends only on observed $p^x_{dj,0}$, the trade cost $\hat{\tau}_{dj}$ from Section D.2, and the preset constants $(\sigma, \mu, k_j/\eta)$ once η is fixed.

Step 2: home cutoff $\hat{a}^{H*}_{j,0}$ from the successful-entry probability. In the model, the probability that a potential entrant survives selection into the home market is the Pareto survival above the home cutoff, $s^H_j = (\underline{a}_j/a^{H*}_{j,0})^{k_j}$. Inverting and using $\underline{a}_j = 1$ from the preset normalization gives

$$\hat{a}^{H*}_{j,0} = \underline{a}_j\,(s^H_j)^{-1/k_j}, \tag{69}$$

which places the home cutoff on the absolute Pareto support without referencing A_j.

Step 3: easiest foreign cutoff $\hat{a}^*_{\text{easy},j,0}$ from the exporter share. The conditional exporter share is the model probability that a home producer also exports to at least one destination, $\pi_j = \Pr(a \geq a^*_{\text{easy},j,0} \mid a \geq a^{H*}_{j,0}) = (a^{H*}_{j,0}/a^*_{\text{easy},j,0})^{k_j}$, the ratio of Pareto survival above the easiest foreign cutoff to survival above the home cutoff. Combining with Step 2,

$$\hat{a}^*_{\text{easy},j,0} = \underline{a}_j\,(s^H_j\pi_j)^{-1/k_j}. \tag{70}$$

This is the foreign cutoff at the destination with the highest draw-survival rate among j's observed export markets.

Step 4: industry productivity scale $\hat{A}_j$. Steps 1 and 3 both characterize the easiest foreign cutoff: Step 1 yields the unit-value-implied $\tilde{a}^*_{\text{easy},j,0} \equiv \min_d \tilde{a}^*_{dj,0}$ measured up to the scale A_j, while Step 3 places it on the absolute support. Equating the two,

$$\hat{A}_j = \frac{\hat{a}^*_{\text{easy},j,0}}{\tilde{a}^*_{\text{easy},j,0}}, \qquad \hat{a}^*_{dj,0} = \hat{A}_j \, \tilde{a}^*_{dj,0}. \tag{71}$$

This identifies A_j as the level shift that aligns the unit-value scale of the cutoffs with their Pareto-survival scale; every remaining foreign cutoff $\hat{a}^*_{dj,0}$ then follows from multiplying its Step 1 counterpart by $\hat{A}_j$.

Step 5: pre-selection mass $\hat{N}_{j,0}$. The model relation $N^H_{j,0} = N_{j,0}\, s^H_j$ says that the observed home establishment stock equals the pre-selection mass times the survival probability. The implied pre-selection mass is therefore the level that scales total baseline revenue in the next step,

$$\hat{N}_{j,0} = \frac{N^H_{j,0}}{s^H_j}. \tag{72}$$

We treat $\hat{N}_{j,0}$ as an intermediate object rather than a structural primitive, since it is mechanically pinned by $(N^H_{j,0}, s^H_j)$ and does not enter the post-shock dynamics separately from β_j and f^E_j.

Step 6: foreign residual-demand shifter $\widehat{\Gamma}_{dj,0}$. The baseline bilateral revenue identity (27) expresses observed export revenue as $R^x_{dj,0} = N_{j,0}\Gamma_{dj,0}(\mu\tau_{dj})^{1-\sigma} A_j^{\sigma-1}\, k_j \underline{a}_j^{k_j} (a^*_{dj,0})^{\vartheta-k_j}/(k_j - \vartheta)$. Solving for $\Gamma_{dj,0}$ and substituting the calibrated objects from Steps 1–5 yields

$$\widehat{\Gamma}_{dj,0} = \frac{R^x_{dj,0}(k_j - \vartheta)}{\hat{N}_{j,0}(\mu\hat{\tau}_{dj})^{1-\sigma} \hat{A}_j^{\sigma-1}\, k_j \underline{a}_j^{k_j} (\hat{a}^*_{dj,0})^{\vartheta-k_j}}. \tag{73}$$

The shifter $\widehat{\Gamma}_{dj,0}$ absorbs the residual destination-market demand needed to rationalize the observed $R^x_{dj,0}$ given the cutoff and the productivity scale.

Step 7: foreign fixed cost $\hat{f}_{dj}$. The foreign zero-profit condition $r_{dj,0}(a^*_{dj,0}) = \sigma f_{dj}$ from the model section expresses the marginal exporter's revenue as $\Gamma_{dj,0}(\mu\tau_{dj})^{1-\sigma} A_j^{\sigma-1}(a^*_{dj,0})^{\vartheta}$. Multiplying through by the Pareto mass above the cutoff and substituting (73) eliminates $\Gamma_{dj,0}$ and gives

$$\hat{f}_{dj} = \frac{R^x_{dj,0}(k_j - \vartheta)\,(\hat{a}^*_{dj,0})^{k_j}}{\sigma\, \hat{N}_{j,0}\, k_j \underline{a}_j^{k_j}}, \tag{74}$$

which expresses the foreign fixed cost as the per-firm revenue at the cutoff divided by σ.

Step 8: home fixed cost $\hat{f}^H_j$. The home zero-profit condition is the analog of Step 7

at the home cutoff $a_{j,0}^{H*}$ with home revenue $R_{j,0}^{H}$, $\tau_{Hj} = 1$, and Pareto mass $(\underline{a}_j / a_{j,0}^{H*})^{k_j}$:

$$\hat{f}_j^H = \frac{R_{j,0}^H (k_j - \vartheta) \, (\hat{a}_{j,0}^{H*})^{k_j}}{\sigma \, \hat{N}_{j,0} \, k_j \underline{a}_j^{k_j}}. \tag{75}$$

Step 9: home expenditure share $\hat{\beta}_j$. The Cobb–Douglas home expenditure identity from the model section reads $\beta_j L_t = R_{j,0}^H$. Evaluated at pre-period aggregate income L_0,

$$\hat{\beta}_j = \frac{R_{j,0}^H}{L_0}. \tag{76}$$

Step 10: sunk entry cost $\hat{f}_j^E$. The free-entry condition from the model section requires that the discounted profit pool from home and foreign markets equals the sunk entry cost. Substituting (74)–(75) into the free-entry algebra yields the closed-form shortcut

$$\hat{f}_j^E = \frac{\vartheta}{\delta_{\text{mo}} \, \sigma \, \hat{N}_{j,0} \, k_j} \left(R_{j,0}^H + \sum_{d \in \mathcal{D}} R_{dj,0}^x \right), \tag{77}$$

in which the entry cost is proportional to total baseline revenue $R_{j,0}^H + \sum_d R_{dj,0}^x$ with a coefficient that depends only on the preset constants, the implied Pareto exponents, and the pre-selection mass.

Equations (68)–(77) recover all the time-invariant fundamentals $\{\hat{A}_j, \widehat{\Gamma}_{dj,0}, \hat{f}_{dj}, \hat{f}_j^H, \hat{\beta}_j, \hat{f}_j^E\}$ needed to evaluate the post-shock equilibrium under the candidate η. The algebra is sequential: each step uses only the data inputs and the quantities produced in earlier steps, so the calibration costs $O(|\mathcal{D}| \cdot |\mathcal{J}|)$ floating-point operations per evaluation of η.

D.4 SMM Implementation

The estimated parameter vector is

$$\theta = (\eta, \alpha, \rho, \nu).$$

For each candidate θ, the baseline objects in Appendix D.3 are recomputed at the candidate value of η. The post-shock equilibrium is then solved using Appendix C, with $M_{Ut} = 1$ and M_{dt} equal to destination d's pre-release AI hiring intensity.

Let m^{emp} collect the empirical moments used in the SMM objective, and let $m^{\text{mod}}(\theta)$

collect the same moments computed from simulated model outcomes. The estimator solves

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \left[m^{\text{mod}}(\theta) - m^{\text{emp}}\right]' W \left[m^{\text{mod}}(\theta) - m^{\text{emp}}\right], \tag{78}$$

where W is diagonal with entries equal to the inverse empirical variances of the targeted reduced-form estimates. The parameter space is

$$\Theta = \{\eta > 0,\ \alpha \geq 0,\ \rho > 0,\ \nu \geq 0\}.$$

The simulated moments use the same product-destination observations, price-group partitions, and destination-AI split as the empirical analysis. They target four reduced-form coefficients:

1. the aggregate HS10 unit-value response;
2. the unit-value response for low-price products shipped to high-AI destinations;
3. the quantity response for the same low-price/high-AI product-destination markets;
4. the aggregate HS10 quantity response.

For each candidate parameter vector, the simulated data contain the pre-period outcome and the post-shock outcome for each relevant observation, and the model moment is the coefficient from the corresponding two-period first-difference projection.

The numerical search is nested because η enters the pre-shock inversion, whereas (α, ρ, ν) govern the post-shock AI counterfactual conditional on that inversion. Given a trial η, the closed-form calibration in Appendix D.3 pins down the baseline primitives: productivity scales, foreign demand shifters, fixed market-access costs, Cobb–Douglas expenditure shares, and sunk entry costs. Holding those primitives fixed, the inner problem searches over (α, ρ, ν) and repeatedly solves the post-shock equilibrium to minimize (78). The outer problem then updates η, rebuilds the baseline calibration, and repeats the inner search. This ordering avoids treating the baseline primitives as fixed objects when the Pareto shape parameter changes.

Inference uses the empirical covariance matrix of the targeted reduced-form moments. For each bootstrap draw, the empirical moment vector is perturbed using the same clustered reduced-form uncertainty, and the full nested SMM problem is re-estimated. The reported standard errors summarize the resulting distribution of parameter estimates.

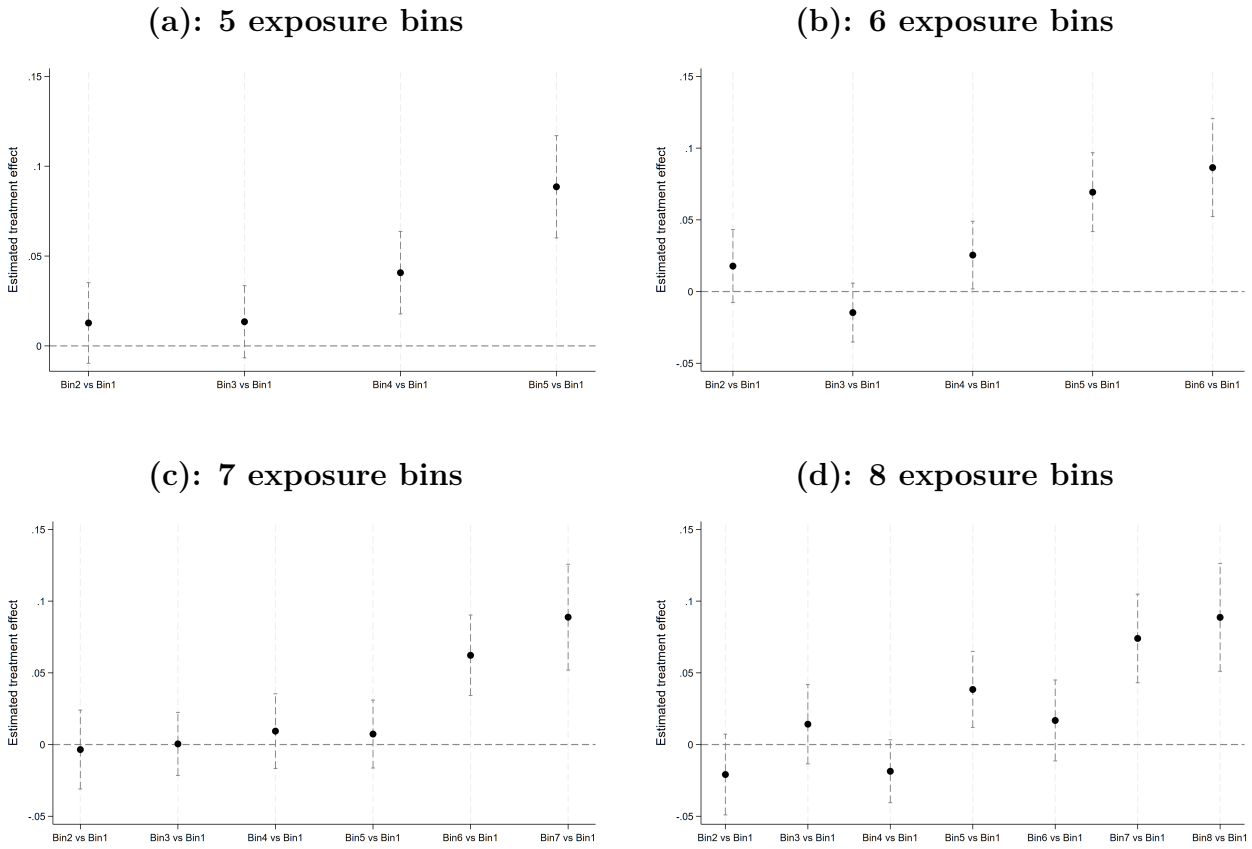
E Tables and Figures

Table A.1: Effects of high GenAI exposure on aggregate export unit values across alternative low-exposure reference groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Bottom 10%		Bottom 15%		Bottom 20%		Bottom 25%	
	Baseline	HS4 trend	Baseline	HS4 trend	Baseline	HS4 trend	Baseline	HS4 trend
GenAIHigh × Post	0.0591***	0.0105	0.0285***	0.0222**	0.0405***	0.0230**	0.0405***	0.0230**
	(0.0114)	(0.0146)	(0.0090)	(0.0110)	(0.0083)	(0.0096)	(0.0083)	(0.0096)
Observations	366,286	366,286	366,286	366,286	366,286	366,286	366,286	366,286
HS10 product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
HS4 × month trend	No	Yes	No	Yes	No	Yes	No	Yes
Clustered SE (HS10)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes. The table reports difference-in-differences estimates of the effect of high GenAI exposure on aggregate export unit value at the HS10 product level from 2020m1 to 2025m12. In each column pair, the high-exposure indicator GenAIHigh_p equals one for HS10 products whose pre-period GenAI exposure exceeds the specified percentile cutoff and equals zero otherwise; products in the bottom 10%, 15%, 20%, or 25% of the pre-period GenAI exposure distribution serve as the reference group for the corresponding column pair. Odd-numbered columns use the baseline specification with HS10 product and year–month fixed effects; even-numbered columns add HS4-specific linear month trends. The dependent variable, Post_t, and fixed effects are defined as in Table 2. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Figure A.1: Estimated effects of GenAI exposure on aggregate export unit values across exposure bins



Notes. The figure displays coefficients of DiD estimation of GenAI-Exposure on Aggregate Export Unit Value by sorting product-level GenAI-Exposure into the number of GenAI-Exposure categories shown in each panel. The estimation is on product level, with time span ranging from 2020m1 to 2025m12. Coefficients are measured relative to the lowest exposure category. Specifications include HS10 product and month fixed effects. Vertical bars show 95% confidence intervals based on standard errors clustered at the HS10 product level.

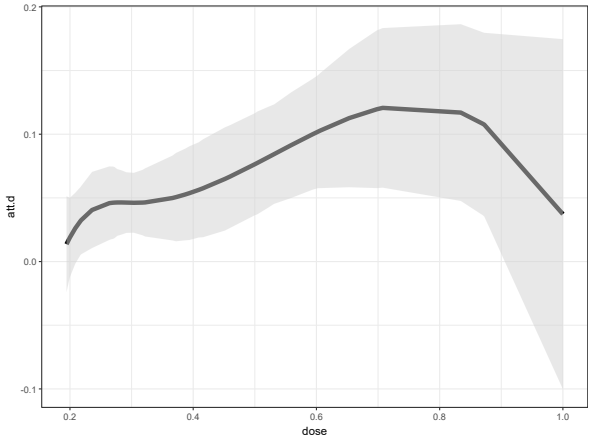
Table A.2: Robustness of the estimated effects of GenAI exposure on aggregate export unit values: continuous treatment intensity

	(1) Bottom 10%	(2) Bottom 15%	(3) Bottom 20%	(4) Bottom 25%
ATT	0.0318***	0.0345***	0.0376***	0.0380***
	(0.0069)	(0.0075)	(0.0078)	(0.0081)
Control share	10.2%	15.6%	20.8%	28.7%
Products	4,127	4,127	4,127	4,127

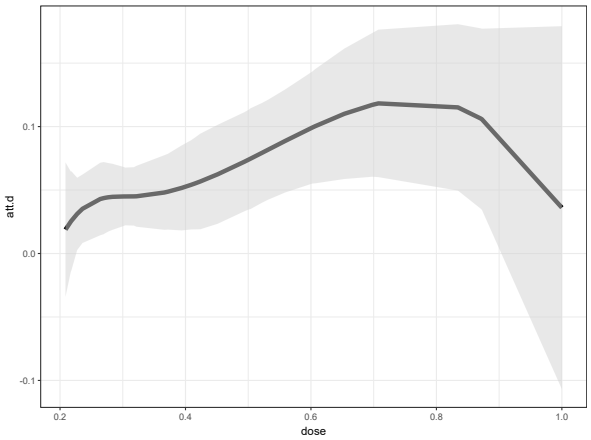
Notes. The table summarizes continuous-treatment difference-in-differences average treatment effects of GenAI exposure on aggregate export unit value on the HS10 product–month panel from 2020m1 to 2025m12, estimated using the nonparametric continuous-treatment difference-in-differences estimator of Callaway, Goodman-Bacon, and Sant'Anna (2024) and implemented through the `contdid` R package. The dose is rescaled GenAI product exposure. In columns (1)–(4), respectively, products in the bottom 10%, 15%, 20%, and 25% of the pre-period exposure distribution serve as the control group. The dependent variable, treatment timing, and fixed effects are defined as in Table 2. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

Figure A.2: Dose–response estimates of the effects of GenAI exposure on aggregate export unit values

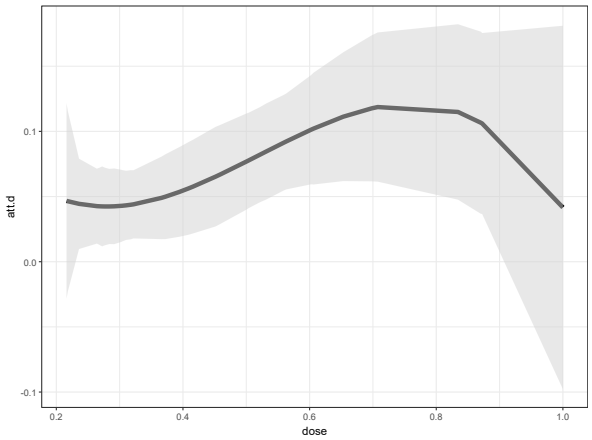
(a): Bottom 10% control



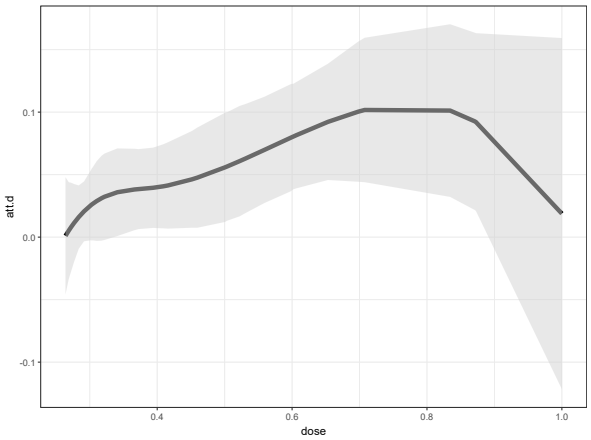
(b): Bottom 15% control



(c): Bottom 20% control



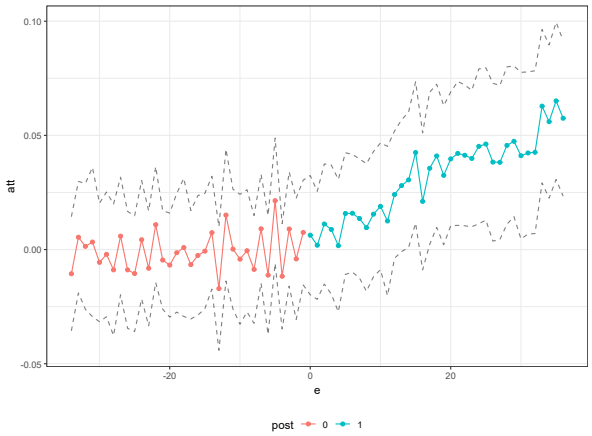
(d): Bottom 25% control



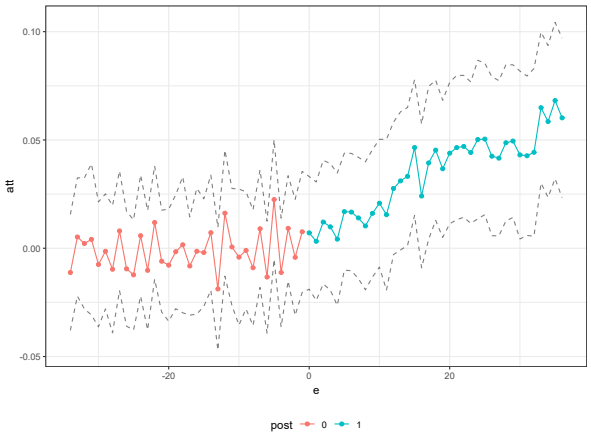
Notes. The figure displays average treatment effects of GenAI exposure on aggregate export unit value on the HS10 product–month panel from 2020m1 to 2025m12, estimated using the continuous-treatment difference-in-differences estimator of Callaway, Goodman-Bacon, and Sant'Anna (2024) and implemented through the `contdid` R package. Panels (a)–(d) use the bottom 10%, 15%, 20%, and 25% of the pre-period GenAI exposure distribution as the control group, respectively. The horizontal axis is the rescaled GenAI dose. Shaded bands show 95% confidence intervals based on standard errors clustered at the HS10 product level.

Figure A.3: Event-study estimates of the effects of GenAI exposure on aggregate export unit values under continuous exposure

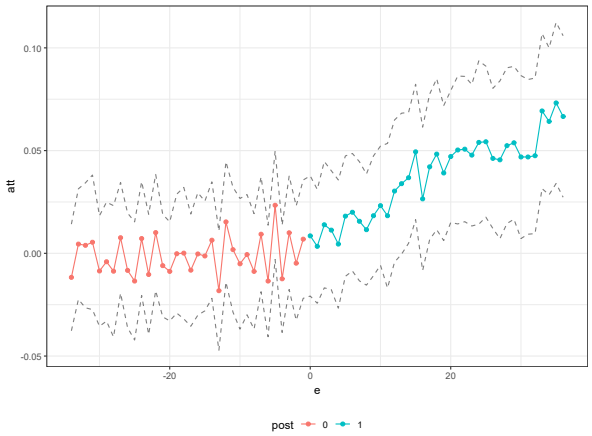
(a): Bottom 10% control



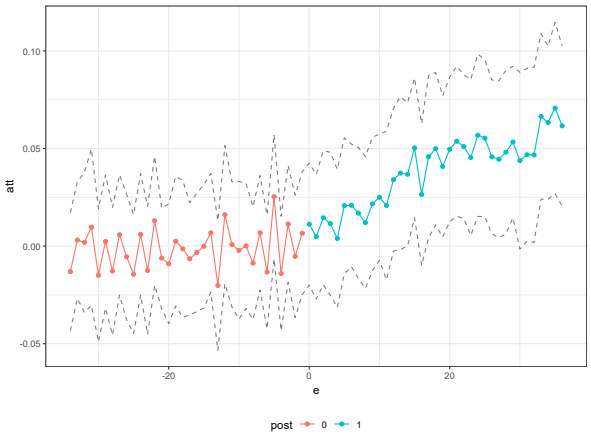
(b): Bottom 15% control



(c): Bottom 20% control



(d): Bottom 25% control



Notes. The figure displays event-study coefficients from the continuous-treatment difference-in-differences estimator of Callaway, Goodman-Bacon, and Sant'Anna (2024), implemented through the `contdid` R package, applied to aggregate export unit value on the HS10 product–month panel from 2020m1 to 2025m12. Panels (a)–(d) use the bottom 10%, 15%, 20%, and 25% of the pre-period GenAI exposure distribution as the control group, respectively. Event time is measured relative to 2022m12, and 2022m11 is the omitted month. Vertical bars show 95% confidence intervals based on standard errors clustered at the HS10 product level.

Table A.3: Alternative estimates of the effects of GenAI exposure on aggregate export unit values: exposure switching

	(1) Bottom 10%	(2) Bottom 15%	(3) Bottom 20%	(4) Bottom 25%
Panel A. Short window				
LATE	0.0259*	0.0286**	0.0374***	0.0374***
	(0.0155)	(0.0125)	(0.0117)	(0.0117)
Observations	125,761	125,761	125,761	125,761
Panel B. Wide window				
LATE	0.0406*	0.0311*	0.0250	0.0250
	(0.0218)	(0.0181)	(0.0164)	(0.0164)
Observations	248,939	248,939	248,939	248,939
Panel C. Full window				
LATE	0.0488*	0.0187	0.0159	0.0159
	(0.0259)	(0.0210)	(0.0189)	(0.0189)
Observations	366,305	366,305	366,305	366,305

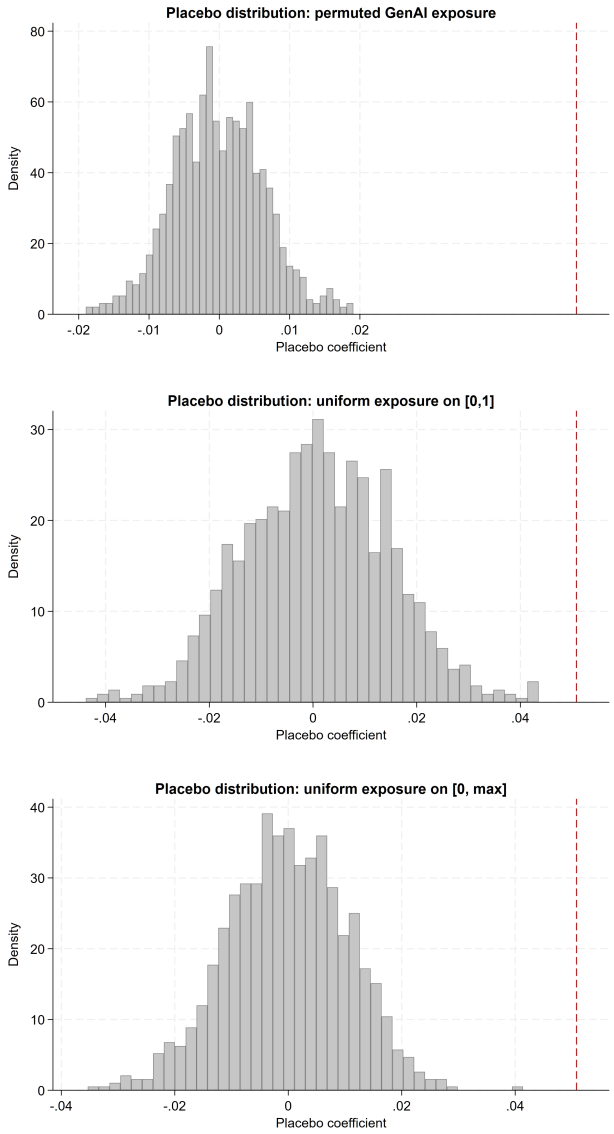
Notes. The table displays fuzzy difference-in-differences local average treatment effects of GenAI exposure on aggregate export unit value through the DiD estimation, conducted on product level. Switchers are products that move above the indicated lower-tail cutoff of the exposure distribution; stayers remain below the cutoff. Panel A uses 2021m12–2023m12, Panel B uses 2020m12–2024m12, and Panel C uses the full 2020m1–2025m12 window. In columns (1)–(4), the cutoff is the bottom 10%, 15%, 20%, and 25% of the exposure distribution, respectively. The outcome, exposure measures, and fixed effects are defined as in Table 2. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table A.4: Effects of GenAI exposure on aggregate export unit values: language-model versus image-generation exposure

	(1)	(2)
	Ln Aggregate Export Unit Value	
	Baseline	With HS4 trend
GenAI × Post	0.0426**	0.0630***
	(0.0167)	(0.0198)
IG-AIIE × Post	0.0069	−0.0190
	(0.0114)	(0.0141)
Observations	366,286	366,286
R^2	0.9639	0.9649
HS10 product FE	Yes	Yes
Month FE	Yes	Yes
HS4 × month trend	No	Yes
Clustered SE (HS10)	Yes	Yes

Notes. The table displays difference-in-differences estimates of the GenAI-Exposure and Image-Generation AI-industry-Exposure on aggregate export unit value, defined as in Table 2 on product level from 2020m1 to 2025m12 . Each column includes $\text{GenAI}_p \times \text{Post}_t$ and $\text{IG-AIIE}_p \times \text{Post}_t$. GenAI_p is pre-release language-model industry exposure; IG-AIIE_p is Image-Generation industry exposure from the same occupational–industry aggregation and HS10 concordance. The regression otherwise follows Table 2. Post_t equals one from 2022m12 onward. Column (1) includes HS10 product and year–month fixed effects; column (2) adds HS4-specific linear month trends. All specifications cluster standard errors at the HS10 product level. Standard errors are in parentheses; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

Figure A.4: Placebo distribution for the estimated effects of GenAI exposure on aggregate export unit values under randomized assignment



Notes. The figure displays the coefficient distribution through the placebo switches, estimating the effect of GenAI-Exposure on Aggregate Unit Value. The time span range from 2020m1 to 2025m12 and the regression is conducted on product level. Each histogram uses 1,000 randomized exposure assignments from the baseline specification in Table A.5. Panel (a) permutes the observed product-level GenAI across HS10 products. Panel (b) draws exposure values from a uniform distribution on $[0, 1]$. Panel (c) draws from a uniform distribution on $[0, \max_p \text{GenAI}_p]$. The dashed red line marks the coefficient estimated with the observed GenAI.

Table A.5: Placebo tests for the estimated effects of GenAI exposure on aggregate export unit values: randomized exposure assignment

	(1)	(2)	(3)	(4)	(5)	(6)
	Permutation		Uniform [0,1]		Uniform [0,max]	
	Baseline	HS4 trend	Baseline	HS4 trend	Baseline	HS4 trend
Actual coefficient	0.0508	0.0404	0.0508	0.0404	0.0508	0.0404
Placebo mean	−0.0001	−0.0002	0.0008	0.0005	−0.0002	−0.0002
Placebo SD	0.0065	0.0065	0.0141	0.0143	0.0107	0.0110
Empirical p	< 0.001	< 0.001	< 0.001	0.008	< 0.001	< 0.001

Notes. The table summarizes randomization placebo tests for the baseline estimate of $\text{GenAI}_p \times \text{Post}_t$ on aggregate export unit value in Table 2. Columns (1)–(2) permute the observed product-level GenAI exposure vector across HS10 products; columns (3)–(4) redraw exposure from a uniform distribution on $[0, 1]$; and columns (5)–(6) redraw exposure from a uniform distribution on $[0, \max_p \text{GenAI}_p]$. Odd-numbered columns use the baseline fixed-effects specification; even-numbered columns add HS4-specific linear month trends. Each placebo distribution uses 1,000 repetitions. The empirical p-value is two-sided and compares the actual coefficient with the placebo distribution.

Table A.6: Sensitivity of the estimated effects of GenAI exposure on aggregate export unit values to industry-level clustering

	(1)	(2)
	Ln Aggregate Export Unit Value	
	Baseline	With HS4 trend
GenAI × Post	0.0508^{***}	0.0404^{***}
	(0.0087)	(0.0093)
Observations	366,286	366,286
R^2	0.9639	0.9649
NAICS clusters	71	71
HS10 product FE	Yes	Yes
Month FE	Yes	Yes
HS4 × month trend	No	Yes
Clustered SE (NAICS)	Yes	Yes

Notes. The table displays difference-in-differences estimates of GenAI-Exposure on aggregate export unit value, replicating Table 2 on industry level from 2020m1 to 2025m12 as a roboustness check. Standard errors are clustered at the four-digit NAICS industry level because GenAI product exposure is assigned through the industry concordance. The dependent variable, sample window, GenAI_p, Post_t, and fixed effects are defined aligning with Table 2. Column (1) is the baseline specification; column (2) adds HS4-specific linear month trends. Standard errors are in parentheses; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.