The Transmission of Trade Shocks in Digital Markets*

Jorge Klinnert†

March 20, 2026

PRELIMINARY DRAFT—PLEASE DO NOT CITE OR CIRCULATE

Abstract

This paper studies how trade policy shocks transmit through digital marketplaces. Using high-frequency product-level data from Amazon, we examine the effects of heightened trade policy uncertainty during the 2025 U.S. presidential transition and the subsequent increase in U.S. tariffs on Chinese imports. We find that trade policy uncertainty reduces the early-life visibility of newly introduced Chinese products, suggesting a worsening of the cold-start problem. The tariff increase leads to incomplete retail price pass-through and substantial adjustment along non-price margins, including fulfillment status, badges, sales rank, sponsored placement, and the composition of brands and manufacturers. To quantify the demand and welfare implications of these changes, we estimate a structural demand model with search position. The model shows that the relationship between visibility and underlying demand strengthens after the tariff shock. Although measured consumer surplus rises in the post-tariff period, it would have increased substantially more absent the contemporaneous change in the platform's ranking rule. The results show that in digital markets, trade shocks propagate through platform-mediated allocation in addition to conventional price channels.

JEL Classification Numbers: L81, E31, F13, F14

Keywords: E-commerce, digital platforms, trade policy, tariffs

*The paper has benefited from conversations with Ginger Zhe Jin, Chenyu Yang, Andrew Sweeting, Nuno Limão, Sharat Ganapati . . . . We acknowledge helpful comments from participants in the Midwest International Trade & Theory Conference. Any remaining errors are our own.

†Department of Economics, University of Maryland. Tydings Hall, College Park, MD 20742, USA. klinnert@umd.edu.

1 Introduction

How do trade shocks transmit when markets are intermediated by digital platforms? A large literature studies the incidence of tariffs through prices, showing how higher import costs are passed through to consumers and firms. But in platform-mediated markets, competition depends on more than price. Platforms determine which products consumers see, how prominently they are displayed, which sellers receive badges, and how information about quality and fulfillment is presented. As a result, when trade costs change, adjustment may occur not only through prices, but also through product visibility, seller composition, and the allocation of consumer attention. Standard pass-through frameworks therefore miss an important feature of digital markets: platforms mediate the transmission of trade shocks.

This paper studies how trade policy shocks affect outcomes in a major digital marketplace. We focus on two closely related events in 2025: a sharp increase in trade policy uncertainty surrounding the U.S. presidential transition and a subsequent escalation of U.S. tariffs on Chinese imports, culminating in a 125 percent tariff increase on April 10, 2025. Using high-frequency product-level data from Amazon search results, we examine how these shocks affect early-life product visibility, prices, non-price attributes, seller composition, search placement, and consumer surplus.

We document two main findings. First, heightened trade policy uncertainty reduces the first-page representation of newly introduced Chinese products relative to comparable non-Chinese products. While these estimates do not directly identify entry, they indicate a deterioration in the early-life visibility of new products, consistent with a worsening of the cold-start problem under heightened uncertainty. Second, the tariff increase leads to incomplete retail price pass-through and substantial adjustment along non-price margins. In more exposed product categories, we observe changes in Prime eligibility, platform-assigned badges, sales rank, sponsored placement, and the composition of brands and manufacturers shown to consumers.

These patterns reflect the fact that digital platforms actively mediate market adjustment. Because consumers typically engage with only the first page of search results, visibility on that page determines effective market access. We show that the mapping from product characteristics to search placement changes after the tariff shock: products with less favorable demand-relevant characteristics become more likely to lose prominence. The platform environment therefore amplifies the competitive consequences of cost shocks by reallocating attention away from products whose sellers increased prices.

To quantify the demand and welfare implications of these changes, we estimate a structural demand model at the

keyword-day level. The model allows price sensitivity and the value of product attributes to differ across the pre- and post-tariff periods and incorporates search position as a determinant of effective demand. Using instruments based on exchange rates, trade policy uncertainty, and differential exposure to Chinese-origin products, we estimate economically meaningful demand elasticities and recover the relationship between visibility and demand. We then compare observed post-tariff consumer surplus with a counterfactual that holds the pre-tariff ranking rule fixed. Measured consumer surplus rises on average in the post-tariff period, but the counterfactual analysis shows that it would have risen substantially more absent the contemporaneous change in ranking. Thus, the platform's ranking adjustment offsets a sizable share of the consumer gains that would otherwise have followed the tariff shock.

This paper contributes to two literatures. First, it extends the literature on tariff incidence by showing that in digital markets, trade shocks affect not only prices and quantities, but also product turnover, visibility, and seller composition, which together shape welfare. Second, it contributes to the literature on digital platforms by showing that platform design and ranking governance influence how macroeconomic shocks transmit through online markets. When platforms allocate consumer attention, the incidence and welfare effects of trade policy depend not only on pass-through, but also on how algorithms map product characteristics into exposure.

More broadly, the paper shows that understanding trade in the digital economy requires integrating cost shocks with platform-mediated competition. As a growing share of retail activity moves online, digital platforms increasingly shape how international trade policy affects consumers, sellers, and market structure.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data, and Section 4 outlines the empirical setting. Section 5 presents the reduced-form methodology and results. Section 6 develops the empirical model and derives its implications. Section 7 concludes.

2 Related Literature

Understanding how trade shocks shape modern retail markets increasingly requires attention to the intersection of international trade and digital platform economics. This paper connects these two strands of literature.

Foundational work in international trade theory emphasizes the distortionary effects of tariffs on resource allocation and welfare. Leith (1971) develops a general equilibrium framework showing how tariffs reshape production and consumption decisions, while Brock and Magee (1978) highlight the political economy forces that shape tariff policy. Melitz (2003) introduces firm heterogeneity into trade models, showing that trade liberalization reallocates resources toward more productive exporters. These frameworks remain central to understanding how

trade policy shocks affect firm-level outcomes and aggregate welfare.

Building on this foundation, recent research highlights the broader and dynamic consequences of tariff shocks. Acemoglu and Tahbaz-Salehi (2025) model how cost shocks propagate through production networks, generating nonlinear effects on firm behavior and input sourcing. Grossman et al. (2024) show that persistent trade frictions can lead to lasting reconfigurations of global supply chains, with implications for market power, product variety, and long-run welfare. Empirical work has largely focused on price transmission: Amiti et al. (2019) find that the 2018 U.S. tariffs were largely absorbed by domestic importers, while Cavallo et al. (2021) and Cavallo et al. (2025) document heterogeneous pass-through across product categories and persistent gaps between import and consumer prices. Related studies show that firms also respond to trade shocks through changes in sourcing, organization, and entry. Fajgelbaum and Khandelwal (2022); Fajgelbaum et al. (2024); Fajgelbaum and Khandelwal (2024) document broad shifts in import patterns and supply chains during the U.S.–China trade war, and Flaaen et al. (2020) and Hiller and Savage (2021) show that tariff hikes disrupted domestic production, profitability, and employment.

A parallel literature in industrial organization and digital economics emphasizes how platform governance shapes competition and market outcomes. Goldfarb and Tucker (2019) argue that platforms act as market designers—rather than neutral intermediaries—by using algorithms that influence search visibility, seller competition, and consumer behavior. Etro (2023) show that platforms adjust access fees and commissions in response to changes in sellers' cost structures, while Teh (2022) develop a framework in which platforms employ non-price governance tools—such as algorithmic visibility, recommendation systems, and entry restrictions—to mediate regulatory shocks. At the firm level, De Loecker (2011) show that multi-product firms respond to trade shocks by adjusting their product mix, which on platforms may translate into withdrawing less profitable items, reallocating marketing efforts, or investing in reputation management to sustain visibility. On the consumer side, Fang et al. (2024) provide experimental evidence that buyer behavior is highly responsive to platform design and information presentation, while Yu (2024) show that advertising and ranking mechanisms influence welfare outcomes in online marketplaces.

This paper bridges these two literatures by examining how trade shocks influence not only prices but also product variety and visibility in digital marketplaces. While the trade literature focuses primarily on price pass-through and the platform literature emphasizes non-price forces that shape competition, we combine these perspectives to show that trade policy can alter the composition of products consumers encounter and choose online. Rather than focusing solely on price adjustments, we examine how trade shocks affect both price and non-price attributes of products in digital marketplaces, and we show how platform design may introduce an additional distortion to

pass-through.

3 Data

Amazon covers about 40 distinct broad product categories (e.g., Electronics, Home & Kitchen, Video Games), each with many narrow product subcategories. We focus on the "Home & Kitchen" category, which includes over 1,300 narrow subcategories (e.g., Towels, Lids, Robotic Vacuums). We focus on the most detailed subcategories in Amazon's category tree to ensure that products within each group are directly comparable (see Figure A.1 in Appendix A.1). To assess the product comparability within each subcategory, we additionally apply a text similarity algorithm to the product titles and descriptions. Figure A.2 and A.3 in Appendix A.1 presents the distribution of similarity scores, along with illustrative examples of products with both low and high similarity values.

We construct a product-level dataset of Amazon search results by using each subcategory name as a keyword. For every keyword, we scrape the first page of search results and record each product's Amazon Standard Identification Number (ASIN), page position, and any platform-provided designations (e.g., "Prime," "Best Seller," "Overall Pick," "Sponsored," "Featured from Amazon Brands"). For each product, we additionally collect its price, average star rating, number of reviews, monthly sales, brand and manufacturer information, title and description, country of origin, and the date the product was first listed.[1] Panel A and Panel B of Figure A.4 in Appendix A.1 provide examples of the interfaces from which we scrape product information.

We repeat this web-scraping process multiple times over a period of slightly less than two months to capture time-dependent variation in search rankings and product attributes. For each product in our sample, we complement these cross-sectional snapshots with high-frequency historical data from a third-party web scraper, including time-varying prices, sales ranks, and referral fees. These data allow us to observe within-product dynamics that a single scraping round cannot capture.

The average product price is $49.81.[2] The average sales rank is 160,996, with an average of 1362 units sold in the previous month. Products in the "Home & Kitchen" category face an average platform commission rate of 14.81%. The average star rating is 4.46, with a mean of 3,370 customer reviews. Finally, 73% of products are Prime. In Appendix A.2, Table A.1 presents summary statistics for the scrapped data.

Panel A of Figure 1 displays the distribution of product origins among items for which country-of-origin

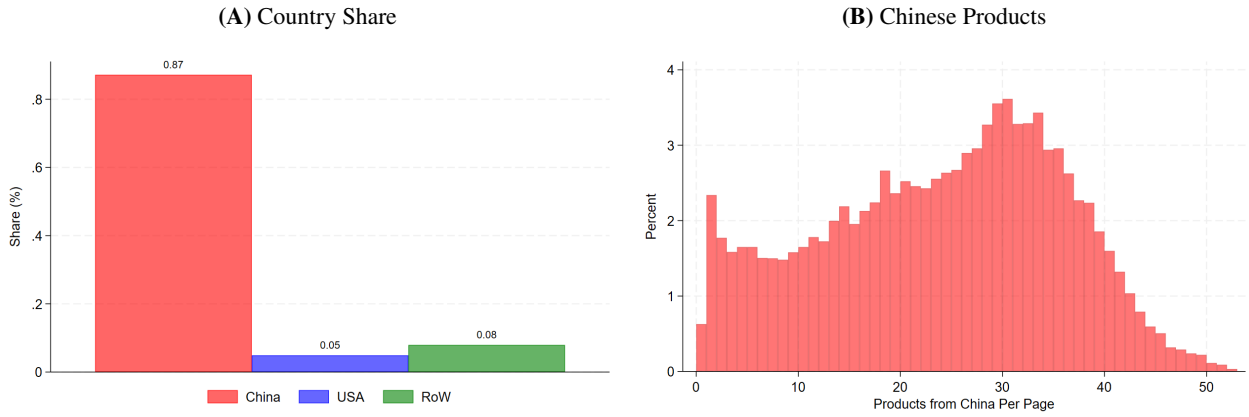
[1] We focus on scraping just the first page of search results, since more than 70% of Amazon shoppers do not click past the first page (FTC, 2023). The first search results page typically displays 60 products.

[2] The average price suggests that, if imported individually, most items would fall below the *de minimis* threshold and would therefore directly affected by its termination.

information is available (88,020 out of 196,474).[3] Approximately 86% of these products originate from China, 6% from the United States, and the remaining 9% from other countries.[4] These figures indicate that at least 86% of products on the platform are potentially affected by recent changes in U.S.–China trade policy. Table A.2 in Appendix A.4 reports the full distribution of products by country, both in levels and percentages.

Panel B of Figure 1 displays the distribution of the average number of China-origin products associated with each keyword. To operationalize variation in exposure to U.S. tariffs on Chinese goods, we classify keywords into "more affected" and "less affected" groups using the 25th percentile of this distribution (14 China-origin products) as a cutoff.

Figure 1. Product Origin



Notes: Panel A presents percentage share of products with origin from China, the U.S., and the rest of the world (RoW) conditional on the county of origin being displayed on Amazon. Panel B presents the distribution of the number of China-origin products per keyword on the first search results page.

4 Empirical Setting

In this section, we outline the empirical setting by discussing the relevant changes to the trade environment.

4.1 Trade Uncertainty

Sudden and substantial changes in trade policy—particularly when announced with little advance notice—can heighten uncertainty for sellers and consumers alike. In the U.S. context, Baker et al. (2016) construct a Trade

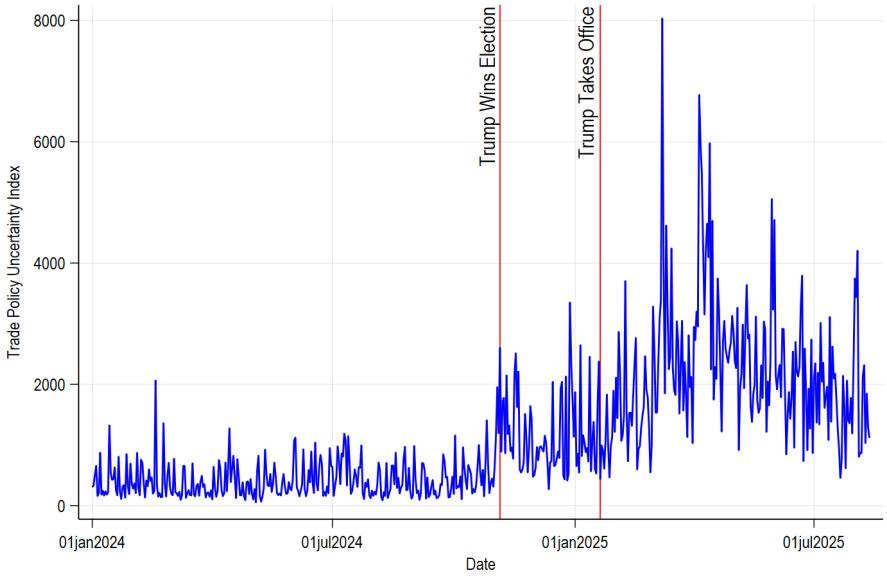
[3] Additional country-of-origin information is currently being extracted using *Rufus*, Amazon's internal large language model. The program accesses Amazon.com and retrieves the product page for each ASIN lacking country-of-origin information. It then submits the query, "*What is the country of origin of this product?*", to *Rufus*. The program records Rufus's response and extracts the country-of-origin information when it is provided.

[4] The predominance of China as the country of origin for products sold on Amazon is consistent with other evidence (see Fleck (2024)).

Policy Uncertainty Index that quantifies this uncertainty using the frequency of articles in leading American newspapers that jointly reference policy-related economic uncertainty and trade policy. The index thus captures periods when trade issues are both salient in the public discourse and perceived as unpredictable in their economic consequences.

Figure 2 plots the daily evolution of the trade policy uncertainty index. The series rises sharply after the November 2024 U.S. presidential election, consistent with the incoming administration's tariff rhetoric and plans to renegotiate trade agreements. A second increase follows the January 2025 inauguration, coinciding with the announcement and implementation of new trade measures. Volatility also intensifies in the aftermath of both events. Taken together, these patterns indicate that the policy environment became simultaneously more restrictive and less predictable, generating heightened uncertainty with possible implications for selling and purchasing behavior.

Figure 2. U.S. Trade Policy Uncertainty Index



Notes: This figure plots the daily U.S. Trade Policy Uncertainty Index measuring the relative frequency of articles in major U.S. newspapers that discuss economic policy uncertainty and make explicit reference to trade policy.

4.2 Tariffs

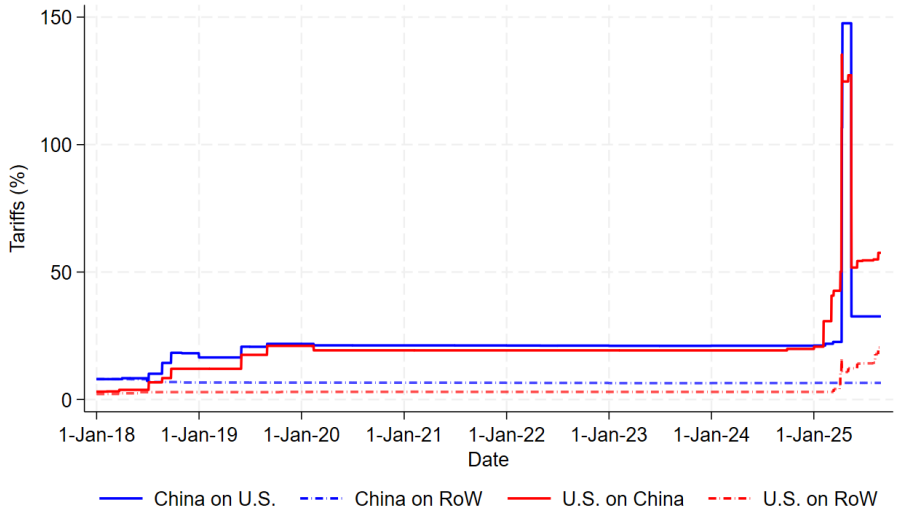
Tariffs remain a central tool of trade policy, allowing governments to influence trade flows, raise revenue, and protect domestic industries. While their use declined in the late twentieth century, recent concerns over supply chain resilience, industrial competitiveness, and economic security have prompted a global resurgence in tariff-based interventions, especially in the U.S..

In this context, U.S. trade policy toward China has entered a new and more aggressive phase. Building on the 2018-2019 trade war, the U.S. maintained elevated tariffs on Chinese imports throughout the early 2020s. In 2025, however, the Trump Administration initiated a sharp escalation. On April 10, 2025, tariffs on Chinese imports were raised to 125%, the most substantial increase since the start of the trade conflict. This action followed the declaration of a national emergency under the International Emergency Economic Powers Act, with the administration citing national security and economic sovereignty as justifications. Figure 6 illustrates the timeline of U.S. tariff actions against China; further details are provided in Appendix A.5 and Appendix A.6.

A related policy change was the termination of the *de minimis* exemption, which had allowed individual shipments under $800 to enter the U.S. duty-free and with minimal customs processing. Widely used by cross-border e-commerce sellers, this exemption had enabled large volumes of low-cost goods from China to reach American consumers with little regulatory oversight. On May 2, 2025, the U.S. revoked *de minimis* eligibility for Chinese and Hong Kong-origin shipments.[5]

Following bilateral negotiations on May 11, the two governments reached a partial resolution. On May 14, 2025, the general tariff on Chinese imports was reduced from 125% to 10%. For sub-$800 shipments, the ad valorem rate was lowered from 120% to 54%, while the $100 flat fee remained in place.

Figure 3. Tariff Trends Timeline



Notes: This figure plots tariff rates between the U.S. and its trading partners—China and the rest of the world (RoW)—as well as between China and its trading partners—the U.S. and the RoW—over time. Data from Bown (2025)

[5]The policy was announced via executive order on April 2, initially imposing a 30% tariff or a $25 flat fee. These were increased to 90% ($75) on April 8 and to 120% ($100) on April 9.

5 Reduced-Form

5.1 Early-Life Visibility

While we are unable to estimate the effect of trade policy uncertainty on overall product entry—because we do not observe products that enter the marketplace but never attain first-page visibility—we can characterize how the representation of newly introduced products among first-page search results evolves around the presidential inauguration that marked a sharp increase in trade policy uncertainty. This is particularly relevant in digital marketplaces, where newly introduced products face a well-documented "cold-start" problem: new products must overcome an initial visibility hurdle to appear in more prominent search positions.

We estimate separate event-study specifications for products originating from China and from non-China countries, relating first-page introductions to calendar time around the inauguration date. Our baseline specification is:

$$y_t = \alpha + \sum_{t \neq t^*} \beta_t \, \mathbb{1}\{\text{Week}_t = t\} + \varepsilon_t, \tag{5.1}$$

where y_t denotes the number of products that are first available in week t and appear on the first page of search results. The indicator $\mathbb{1}\{\text{Week}_t = t\}$ equals one for observations corresponding to calendar week t, and t^* denotes the omitted reference week of the presidential inauguration.

In Equation 5.1, the parameters of interest, β_t, capture deviations in the representation of newly introduced products among first-page search results in week t relative to the pre-inauguration baseline. These estimates reflect changes in early-life visibility conditional on entry, rather than changes in product entry per se.

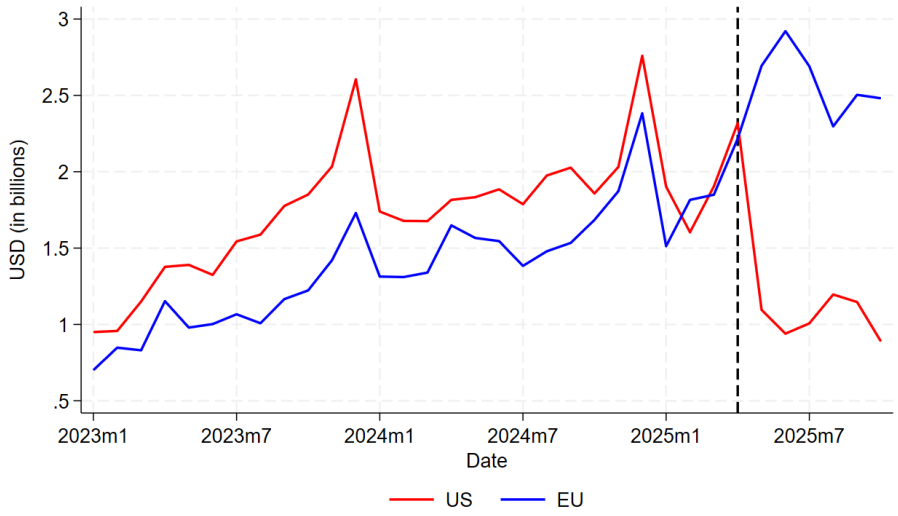
The underlying economic mechanism is that the presidential inauguration introduced substantial uncertainty regarding future trade policy, tariff schedules, and market access conditions. In response to this heightened uncertainty, sellers may delay the introduction of new products until the policy environment becomes clearer. Such delays translate into a discrete reduction in the number of newly introduced products that are able to overcome the cold-start constraint and attain first-page visibility.

5.1.1 Results

We begin by using Chinese export data on low-value shipments, which are more likely to be sold on Amazon and also more exposed to the de minimis exemption. We focus on exports to the United States, which was directly affected by the trade policy changes, and to the European Union, which was not. Figure 4 plots the resulting time

series before and after the implementation of the tariffs. Exports to the European Union continue to follow a stable trend, while exports to the United States exhibit a marked decline following the policy change, consistent with increased trade uncertainty and higher policy-induced costs.

Figure 4. Chinese Low-Value Exports



Notes: This figure shows exports recorded under Chinese customs codes 9804, 9805, and 98010010 (General Administration of Customs of the People's Republic of China, 2026). The red line denotes exports to the United States, and the blue line denotes exports to the European Union. Export values are measured in billions of U.S. dollars.

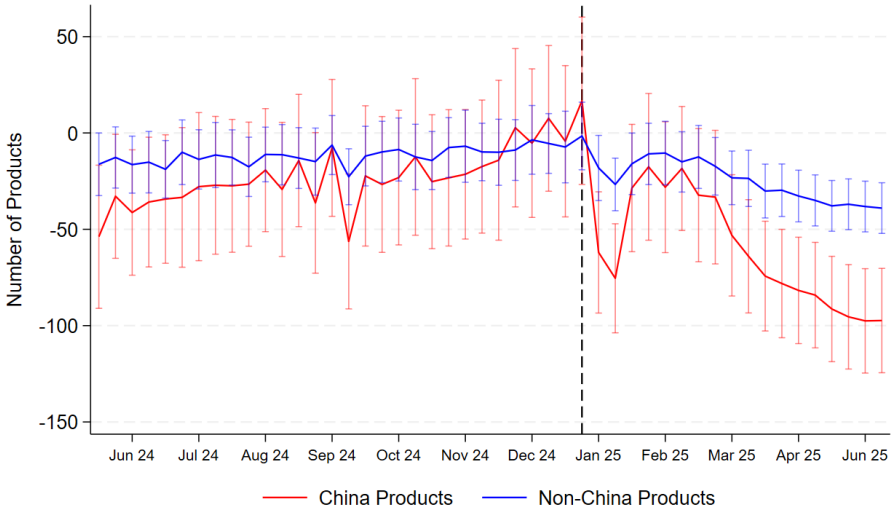
Next, we examine the representation of newly introduced products among first-page search results before and after the increase in trade policy uncertainty associated with the 2025 presidential transition, using the specification in Equation 5.1. Figure 5 reports the estimated coefficients separately for products originating in China and products originating elsewhere. Before the inauguration, both series are relatively stable, with no clear evidence of differential pre-trends between China and non-China products. This suggests similar pre-period dynamics in early-life visibility across the two groups.

After the inauguration, the two series diverge sharply. Newly introduced products from China experience a large and persistent decline in first-page representation, with coefficients becoming substantially more negative in the weeks immediately following the event and continuing to fall thereafter. By contrast, newly introduced products from non-China countries show only a modest decline, with estimates remaining much closer to zero throughout the post-inauguration period.

Taken together, these findings point to a contraction in the flow of low-value Chinese goods to the United States and a concurrent decline in the early-life visibility of newly introduced Chinese products on Amazon. While the

visibility estimates do not identify changes in entry directly, they indicate that new Chinese products become less likely to secure first-page placement during this period. This pattern thus suggests a worsening of the cold-start problem under heightened trade policy uncertainty.

Figure 5. Early-Life Visibility



Notes: the figure depicts event-study estimates of the number of newly introduced products appearing on the first page of search results around the 2025 presidential inauguration. The red line denotes products originating from China and the blue line denotes products from non-China countries. Coefficients are relative to the week preceding the inauguration; vertical bars indicate 95 percent confidence intervals.

5.2 Prices and Product Characteristics

To estimate the causal effect of tariffs on product outcomes, we employ a difference-in-differences (DiD) framework that exploits the sudden increase in U.S. tariffs on Chinese imports implemented on April 10, 2025. We construct subcategory-level panels by aggregating product-level data for items appearing on the first page of search results. The identification strategy relies on cross-sectional variation in tariff exposure across product subcategories, defined by their pre-policy share of imports originating from China. Our main estimating equation is given by:

$$y_{ct} = \alpha + \beta \textit{High-Exposure}_c + \gamma \textit{Post}_t + \delta(\textit{High-Exposure}_c \times \textit{Post}_t) + \kappa_c + \varepsilon_{ct}. \tag{5.2}$$

In this specification, $y_{ct} = \sum_{i=1}^{n} \frac{y_{it}}{n}$ denotes the outcome variable of interest for subcategory c in period t.[6] The variable $\textit{High-Exposure}_c$ is a binary indicator equal to 1 if subcategory c contains more than 14 products originating

[6] y_{it} denotes the outcome variable for product i at time t. Subcategory c has $1, \ldots, n$ products.

from China, and 0 otherwise.[7] Similarly, $Post_t$ is a binary indicator equal to 1 for periods after April 10, 2025, and 0 otherwise. We also include keyword fixed effects, denoted by κ_c, to control for time-invariant differences across search terms.

By running the regression at the subcategory level, we control for competition between Chinese and non-Chinese sellers within the same narrowly defined product markets. This approach allows us to compare subcategories that are broadly similar in terms of product type and demand characteristics, but that differ in their exposure to Chinese imports. In doing so, we are able to isolate the impact of the tariff increase on subcategories with Chinese products by comparing them to subcategories with no Chinese-origin products. This design helps ensure that any observed effects are driven by differential exposure to the tariff shock—rather than by unrelated differences across broader product categories. Nevertheless, we additionally estimate the model at the product level, allowing us to incorporate product fixed effects and to test whether our subcategory-level results mask within-market variation.[8]

In Equation 5.2, the parameter of interest, δ, measures the causal effect of the abrupt increase in tariffs on Chinese goods across a range of subcategory-level outcomes. The primary economic mechanism underlying this relationship is tariff pass-through: when sellers face higher import costs due to the policy and do not fully absorb these costs, part of the burden is transmitted to consumers through higher prices. Higher prices may, in turn, lead the platform to partially absorb the tariff shock through adjustments visibility to limit demand responses. The residual incidence borne by consumers is then reflected in observed changes in consumer behavior.

5.2.1 Results

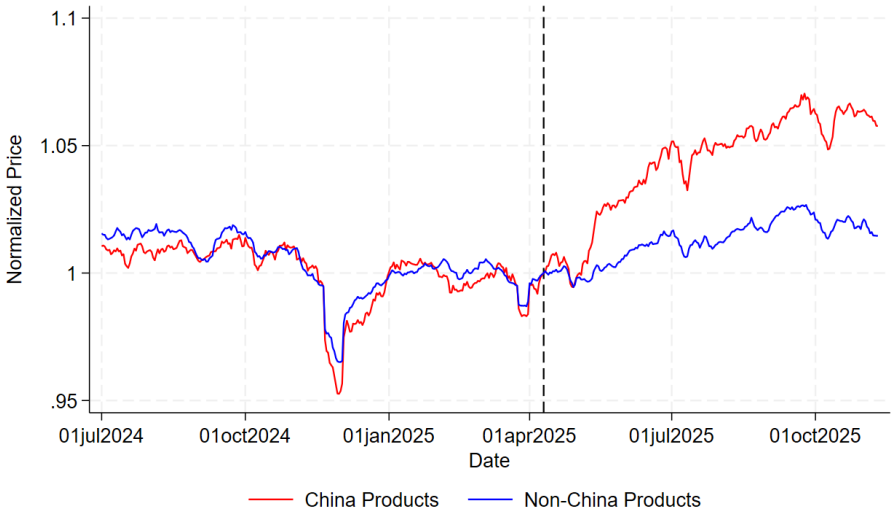
We begin by using the high-frequency price data together with information on each product's country of origin to evaluate the impact of the tariffs on prices. We classify a product as affected if its country of origin is China and unaffected otherwise. For both groups, we normalize prices by their level on April 10, 2025, which serves as the pre-policy baseline. Figure 6 plots the resulting time series of normalized prices before and after the implementation of the tariffs. While prices of non-Chinese products exhibit a modest increase following the tariff announcement, Chinese products experience a substantially larger rise.

We continue by examining how tariffs affect product characteristics on the platform using our scraped data. Table A.6 summarizes the results.

[7] We treat products originating from Hong Kong as Chinese.

[8] Establishing parallel trends is crucial to ensure that post-tariff shock differences can be interpreted causally as effects of the shock rather than reflecting pre-existing trends. We therefore test the parallel trends assumption by examining pre-shock differences in our main variables between China and non-China products across generations before the tariff shock. The results of this pre-trend analysis are reported in Table A.4 of Appendix A.8.

Figure 6. Price Time-Series



Notes: This figure plots average normalized prices for products classified as affected (China-origin) and unaffected (non–China-origin) by the U.S. tariffs. Prices for each group are normalized to their level on April 10, 2025. The dashed vertical line marks the implementation date of the tariffs.

Consistent with partial pass-through, tariffs led to a modest but statistically significant increase in prices in more exposed subcategories. Column (1) shows that the interaction between high exposure and the post period is positive and precisely estimated, implying that prices rose by roughly 1.8 percent more in high-exposure subcategories than in less exposed subcategories after the tariff shock. At the same time, Prime eligibility declines in exposed markets (Column 2), pointing to some deterioration in fulfillment conditions following the policy change. We also find a small decline in the share of products carrying the "Best Seller" badge (Column 3), while there is no detectable effect on the likelihood of receiving the "Overall Pick" badge (Column 4).

Sales performance deteriorates sharply in the treated subcategories. Column (5) shows that more exposed subcategories experience a statistically significant worsening in sales rank after the tariff, consistent with lower realized demand or weaker sales performance relative to less exposed markets. By contrast, we do not find meaningful changes in average ratings or in the number of comments (Columns 6 and 7), suggesting that the tariff shock does not materially affect these measures of customer engagement in the short run.

The tariff shock also changes the composition of products shown on the platform. Columns (8) and (9) show significant increases in both brand variety and manufacturer variety in high-exposure subcategories after the tariff. Rather than reflecting a genuine expansion in choice, these increases are more naturally interpreted as evidence of greater turnover and relisting as sellers adjust their product mix in response to the change in trade costs.

Finally, Column (10) shows that the number of sponsored products increases in high-exposure subcategories after the tariff. This pattern is consistent with an increase in sellers' demand for paid visibility in the markets most affected by the shock and with the platform accommodating that demand through greater advertising placement. In this sense, adjustment to the tariff occurs not only through prices and assortment, but also through the platform's advertising margin.

Overall, the evidence shows that tariffs affect several product characteristics on the platform, not only prices. This broader set of responses suggests that the sellers and the platform adjust along multiple margins following the increase in trade costs.

Table 1. Tariff Results

	Price	Prime Delivery	Best Seller	Overall Pick	Sales Rank	Ratings	Com-ments	Brand Variety	Manuf. Variety	# Sponsored
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High-Exposure	-0.004	0.054***	-0.007***	-0.004***	0.095***	-0.001	-0.020	0.167***	0.162***	0.196***
	(0.010)	(0.012)	(0.002)	(0.001)	(0.027)	(0.001)	(0.038)	(0.022)	(0.022)	(0.026)
Post	-0.030***	-0.353***	0.002*	-0.001*	-0.025*	-0.001*	-0.369***	-0.009	-0.014*	-0.002
	(0.006)	(0.006)	(0.001)	(0.000)	(0.015)	(0.000)	(0.019)	(0.008)	(0.008)	(0.010)
High-Exposure × *Post*	0.018***	-0.029***	-0.002**	0.000	0.049***	-0.000	-0.013	0.029***	0.032***	0.019*
	(0.007)	(0.007)	(0.001)	(0.000)	(0.017)	(0.000)	(0.022)	(0.008)	(0.008)	(0.011)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782

Notes: This table shows the estimates from Equation 5.2. Columns (1) shows results related to price, Columns (2)-(10) show results related to delivery, platform assigned badges, sales, ratings, comments, brand variety, manufacturer variety, and number of sponsored products. Robust standard errors are clustered at the ASIN-level and in parenthesis. Statistical significance is denoted by * p<0.1, ** p<0.05, and *** p<0.01.

In Appendices A.9 and A.10, we conduct a series of robustness checks showing that these results remain consistent when using an alternative timing of the exogenous shock and when estimating the model at different levels of disaggregation.

5.3 Product Placement in Search Results

The platform determines each product's placement through a ranking algorithm that maps observable product features into a placement index. Since the exact algorithm is opaque, we approximate this relationship as follows:

$$r_{jct} = \mathbf{X}'_{jct}\beta_x + \zeta_j + \kappa_c + \tau_t + \epsilon_{jct}, \tag{5.3}$$

where $\mathbf{X}_{jct}$ denotes the vector of product attributes—such as price, badges, ratings, and sales rank—that enter the platform's ranking algorithm; ζ_j captures time-invariant product quality; κ_c absorbs keyword-level normalization across search markets; and τ_t accounts for time-specific shifts in the platform's ranking environment. The residual ϵ_{jct} represents unobserved, idiosyncratic variation in placement.

To evaluate how the platform's ranking algorithm evolved around the tariff event, we estimate Equation 5.3 separately for the pre- and post-tariff periods, denoting by β_z^0 and β_z^1 the implicit algorithmic weights on observable features before and after the tariff, respectively. Using these estimated coefficients, we predict each product's counterfactual placement under four combinations of algorithms and attributes: the pre-tariff algorithm applied to pre-tariff attributes (actual outcome), the pre-tariff algorithm applied to post-tariff attributes, the post-tariff algorithm applied to pre-tariff attributes, and the post-tariff algorithm applied to post-tariff attributes (actual outcome). These counterfactual predictions allow us to decompose the overall change in placement between the two periods into two distinct components:

$$\underbrace{\widehat{r}^{\,1,1} - \widehat{r}^{\,0,0}}_{\text{Total Change}} = \underbrace{\left(\widehat{r}^{\,0,1} - \widehat{r}^{\,0,0}\right)}_{\text{Attribute Change}} + \underbrace{\left(\widehat{r}^{\,1,1} - \widehat{r}^{\,0,1}\right)}_{\text{Algorithmic Change}} . \tag{5.4}$$

The first term on the right-hand side captures how changes in observable attributes—such as price adjustments, ratings, or badge assignments—would have affected placement if the platform's ranking rule had remained fixed. The second term isolates how the platform's algorithm reweighted these same attributes over time, holding product characteristics constant. The sum of the two components equals the total observed change in product placement. A positive value (for any component) indicates that the change led to products being placed lower on the search page—reflecting worse visibility—whereas a negative value implies that the change improved visibility by moving products higher on the page.

5.3.1 Results

Table 2 reports estimates of the relationship between product placement and observable product attributes before and after the implementation of tariffs. The coefficients can be interpreted as the implicit algorithmic weights the platform assigns to each attribute when determining search rankings. The estimated coefficient of price increases from 4.5 to 5.9, indicating that an increase in price is associated with a worse rank position after the tariff, holding other attributes constant. This implies that the platform's algorithm became more price-sensitive following the tariff event, which would dampen firms' ability to fully pass through higher costs into prices, reinforcing the mechanism of incomplete pass-through discussed earlier.

Table 2. Algorithmic Sensitivity of Product Visibility

	Price	Prime	Best Seller	Overall Pick	Rating	Sales Rank
	(1)	(2)	(3)	(4)	(5)	(6)
Position Pre	4.503***	-1.497***	0.576**	-9.522***	-1.772***	4.071***
	(0.283)	(0.130)	(0.250)	(0.308)	(0.339)	(0.097)
Position Post	5.933***	-0.823***	0.574*	-10.494***	-1.246***	4.100***
	(0.400)	(0.112)	(0.301)	(0.348)	(0.399)	(0.113)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations Pre	535,596	535,596	535,596	535,596	535,596	535,596
Observations Post	342,307	342,307	342,307	342,307	342,307	342,307

Notes: This table reports estimates of the relationship between product position and product attributes before and after the implementation of tariffs. "Pre Position (0)" corresponds to the pre-tariff attribute estimates, $\widehat{\zeta}_x^0$, and "Post Position (1)" corresponds to the post-tariff estimates, $\widehat{\zeta}_x^1$. All specifications include keyword, product (ASIN), and time fixed effects. Standard errors are clustered at the keyword–ASIN level and shown in parentheses. Statistical significance is denoted by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Using the pre- and post-tariff models to predict each product's placement under alternative combinations of algorithm and product attributes, allows us to decompose observed changes in placement into an attribute-driven component and an algorithmic component.We therefore quantify how much of the change in visibility reflects the direct effect of evolving product characteristics and how much reflects the platform's re-optimization of its ranking rule.

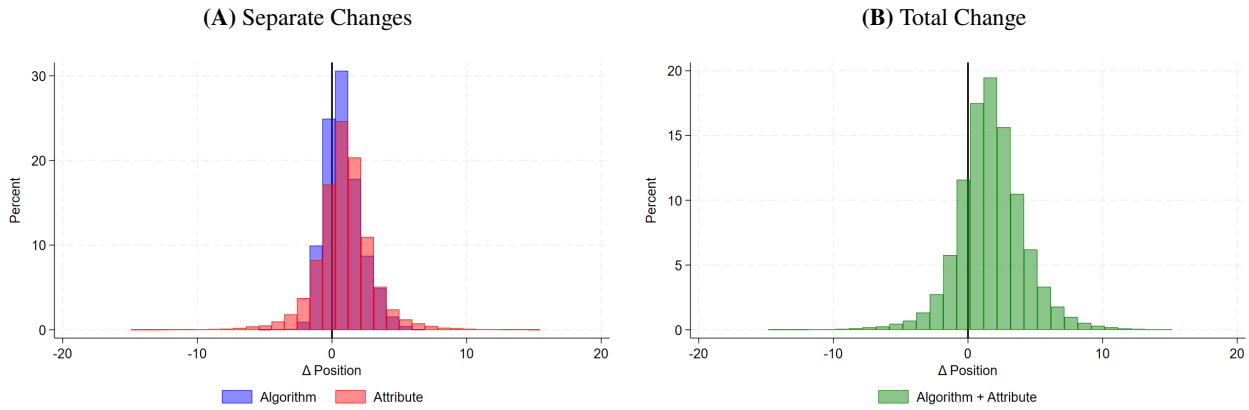
Figure 7 presents the distributions of the attribute-driven effect, the algorithmic effect, and their sum. Panel A shows that the attribute-driven component (green) is shifted slightly to the right of zero, indicating that changes in observable product attributes alone would have led products to be placed lower on the search page following the tariff increase. The algorithmic component (blue) is also right-shifted, implying that the platform's post-tariff ranking algorithm further penalizes these products by assigning less favorable placements, holding attributes constant.

Panel B displays the distribution of the total effect (red), which reflects the combined impact of changes in product attributes and algorithmic reweighting. The distribution is shifted further to the right, indicating that these two forces reinforce one another: attribute-level adjustments reduce visibility, and contemporaneous changes in the ranking algorithm amplify this effect.

This pattern is consistent with a platform strategy aimed at preserving consumer demand during a period of elevated tariffs. Rather than allowing tariff-induced cost increases to fully translate into a worse consumer experience, Amazon appears to have adjusted visibility in ways that helped keep shoppers on the platform, even

if doing so shifted a larger share of the burden onto affected sellers. In this sense, the evidence is consistent with Amazon competing aggressively to retain consumer traffic during the tariff episode (Nassauer and Khan, 2025). It is also broadly aligned with the firm's stated emphasis on customer obsession.[9]

Figure 7. Changes in Product Placement



Notes: The figure decomposes the change in predicted product placement after the tariff into an attribute-driven component and an algorithmic component. Panel A shows their separate distributions; Panel B shows the distribution of the total effect.

6 Structural Model

We assume the following timing. First, the seller sets the price of the product. Second, the platform observes prices and product characteristics and determines each product's search position. Finally, consumers observe the price, search position, and other characteristics of products displayed for a given keyword on a given day and choose whether and which product to purchase.

6.1 Consumer

Our reduced-form results demonstrate that trade shocks in digital marketplaces transmit through multiple channels beyond prices. To quantify the demand and welfare implications of these channels, we estimate a structural demand model.

We implement a simple logit model in which a market m is defined as a search keyword–day pair. Within each market, consumers choose among products displayed on the first search results page and an outside option representing not purchasing any product from that page. Let j index products and let $m = (k, t)$ denote the

[9] Although Amazon lists 16 Leadership Principles, its best-known guiding principle is customer obsession.

keyword–day market. The indirect utility that consumer i derives from product j in market m is given by

$$u_{ijm} = \begin{cases} -\alpha_0 p_{jm} + \beta_0 r_{jm} + \mathbf{X}'_{jm}\gamma_0 + \mu_j + \lambda_k + \xi_{jm} + \varepsilon_{ijm}, & \text{Post} = 0, \\ -\alpha_1 p_{jm} + \beta_1 r_{jm} + \mathbf{X}'_{jm}\gamma_1 + \mu_j + \lambda_k + \xi_{jm} + \varepsilon_{ijm}, & \text{Post} = 1. \end{cases} \tag{6.1}$$

where we allow preference parameters to differ between the pre- and post-tariff periods to accommodate potential changes in consumer demand following the tariff shock.[10] Moreover, in this specification, p_{jm} denotes price, r_{jm} denotes position, and X_{jm} is a vector of observed product characteristics. The terms μ_j and λ_m represent ASIN and keyword fixed effects, respectively. The component ξ_{jm} captures time-varying unobserved product-market quality, and ε_{ijm} is assumed to follow a Type I extreme value distribution. Under the logit assumption, the market share of product j relative to the outside option in regime τ, where $\tau \in \{0, 1\}$ indexes the pre- and post-tariff periods, satisfies

$$\ln s_{jm} - \ln s_{0m} = -\alpha_\tau p_{jm} + \beta_\tau r_{jm} + \mathbf{X}'_{jm}\gamma_\tau + \mu_j + \lambda_k + \xi_{jm}. \tag{6.2}$$

To implement the model in the data, we introduce several assumptions. First, because the dataset reports sales ranks rather than quantities, we convert sales ranks into sales using the mapping procedure from He and Hollenbeck (2020). Second, because sales ranks and thus the implied sales quantities are observed at the platform level rather than at the keyword level, we allocate product-level sales across keywords. Specifically, for each product, we compute the share of total search volume attributable to the focal keyword relative to the total search volume across all keywords in which the product appears, and use this ratio to apportion platform-level sales to the keyword level. Finally, following Yu (2024) we calibrate market size using industry-level benchmarks. Taken together, these steps allow us to construct keyword–day product-level market shares. Appendix A.11 provides details on aforementioned steps.

Price may be endogenous if it is correlated with unobserved product–market quality, ξ_{jm}. To identify the price sensitivity parameter α_τ, we instrument for price using interactions between fluctuations in the U.S.–China exchange rate, the Daily Trade Policy Uncertainty Index, and the number of products in a keyword market originating from China. Because a substantial share of sellers are based in China, these interactions plausibly shift marginal costs through exchange rate and policy exposure. At the same time, conditional on product and market fixed effects, they are unlikely to directly affect U.S. consumer demand except through prices. This strategy generates exogenous

[10] Tariffs may have altered price sensitivity or the valuation of observable product attributes, including badges, ratings, comments, and position. Estimating tariff period-specific coefficients permits the demand system to adjust across regimes rather than imposing a common preference structure throughout the sample.

price variation within keyword–day markets.

6.1.1 Results

We estimate the demand equation, Equation 6.2, using instrumental variables with keyword and product fixed effects and standard errors clustered at the product level.

Table A.9 reports the demand parameter estimates, and Figure 8 visualizes some key results. Consistent with price endogeneity, the OLS estimates understate the magnitude of price sensitivity, while the IV estimates in Table A.9 yield larger (in absolute value) and economically meaningful coefficients. Appendix A.11.3 reports first-stage results for the instruments used to identify price sensitivity. Panel A of Figure 8 reports the distribution of product-level own-price elasticities implied by the estimated demand model. We find that implied elasticities generally range between 3 and 5 in absolute value, consistent with elastic demand in online retail markets (Gutierrez, 2021; Farronato et al., 2023; Yu, 2024). Moreover, demand becomes more price elastic in the post-tariff period, indicating increased consumer responsiveness to price changes following the trade shock.

Accounting for search position is essential. Even if two products have identical underlying utility, the product ranked higher on the results page is more likely to be viewed and thus captures a larger market share. We therefore include search position as a control in the demand specification.

To study how underlying quality varies across positions—and how this relationship changes after the tariff—we construct the figures separately for the pre- and post-tariff periods. Using the estimated mean utility, we normalize product-level utility to have mean zero and unit variance within each market and then compute the average normalized utility by position. This allows a direct comparison of the relationship between rank and underlying demand across periods. In both periods, higher-ranked products exhibit higher normalized utility, with average utility declining as position worsens. Post tariff, however, this gradient becomes steeper, indicating a stronger association between rank and underlying demand after the shock.

We next examine heterogeneity by platform-assigned badges. Products carrying either the "Best Seller" or "Overall Pick" badge display consistently higher normalized utility than non-badged products at comparable positions in both periods. This separation becomes steeper post tariff, implying that the association between badges and underlying demand strengthens after the policy change. At the same time, the two badges differ in their distribution across the first page: "Best Seller" products appear throughout the page, whereas "Overall Pick" products remain concentrated near the top positions.

Finally, we consider consumer feedback signals. Differences in normalized utility across high versus low

star ratings are relatively modest in both periods. By contrast, products with a larger number of reviews exhibit systematically higher normalized utility, especially at lower-ranked positions. This relationship also becomes steeper post tariff, suggesting that review volume becomes more strongly associated with underlying demand after the tariff shock. Taken together, these patterns indicate that, following the tariff, rank becomes more tightly aligned not only with estimated demand overall, but also with badge status and review volume.

6.2 Seller

We model sellers as single-product firms that set prices to maximize static profits, taking the platform's ranking rule as given. In market m, seller j earns profits

$$\pi_{jm}(p_{jm}) = \left(p_{jm} - mc_{jm}\right) M_m s_{jm}, \tag{6.3}$$

where M_m denotes market size, s_{jm} is the demand share implied by (6.2), and mc_{jm} denotes marginal cost. The seller chooses p_{jm} to maximize π_{jm}.

An important feature of our model is that the platform assigns search positions after observing prices and other product characteristics. Consequently, price affects demand through two channels: *(i)* directly, via its effect on consumer utility, and *(ii)* indirectly, through its effect on product visibility. In regime τ, search position is determined by the ranking rule $r_{jm} = f_\tau(p_{jm}, \mathbf{X}_{jm})$. Let α_τ denote the marginal utility of price and β_τ the coefficient on position in mean utility. Differentiating mean utility with respect to price yields

$$\frac{d\delta_{jm}}{dp_{jm}} = -\alpha_\tau + \beta_\tau \frac{\partial f_\tau(\cdot)}{\partial p_{jm}}. \tag{6.4}$$

Under the logit demand system, the own-price elasticity of product j in market m under regime τ is

$$\varepsilon_{jjm} = \frac{p_{jm}}{s_{jm}} \frac{ds_{jm}}{dp_{jm}} = p_{jm}(1 - s_{jm}) \left(-\alpha_\tau + \beta_\tau \frac{\partial f_\tau(\cdot)}{\partial p_{jm}}\right). \tag{6.5}$$

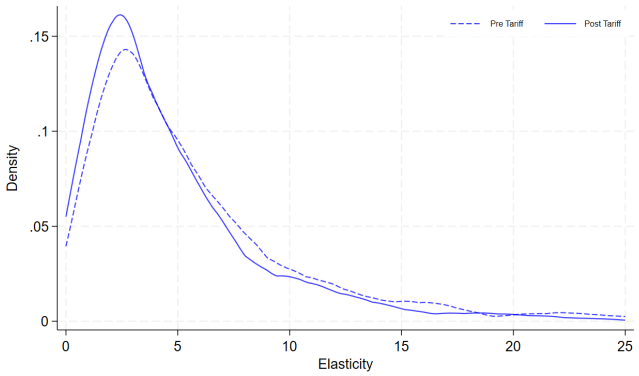
The seller's first-order condition can therefore be written as

$$\frac{p_{jm} - mc_{jm}}{p_{jm}} = \frac{1}{|\varepsilon_{jjm}|}. \tag{6.6}$$

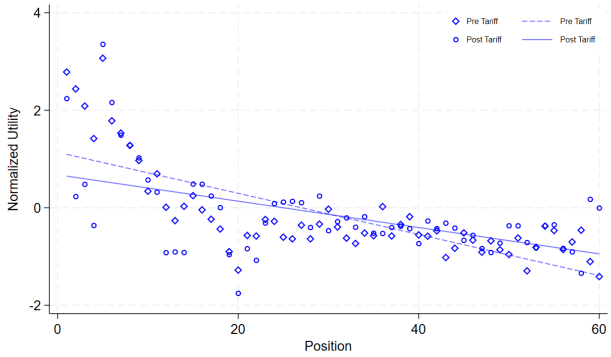
Combining observed prices with the estimated elasticities, we use this relationship to recover the implied markup

Figure 8. Consumer Estimation Results

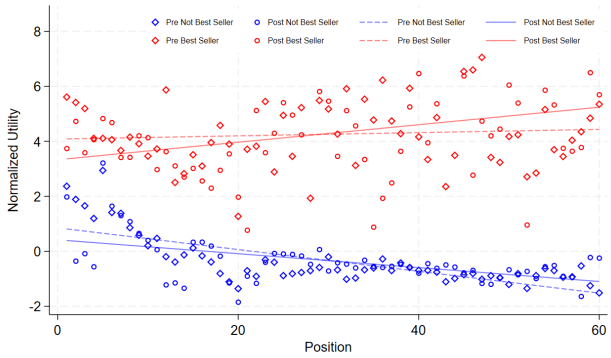
(A) Price Elasticity



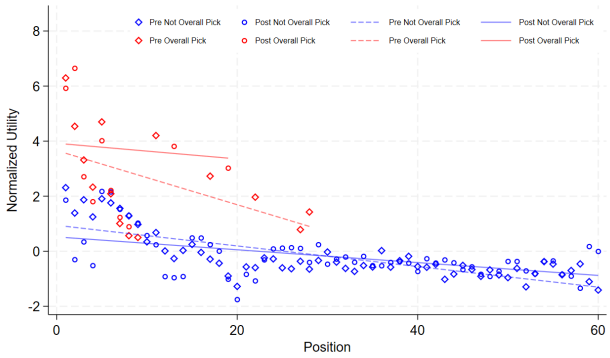
(B) Product Utility vs. Position



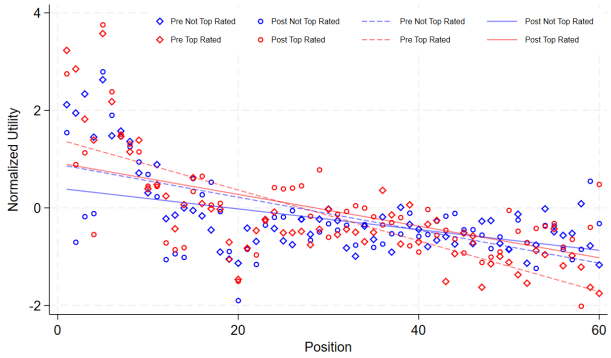
(C) Product Utility vs. Best Seller



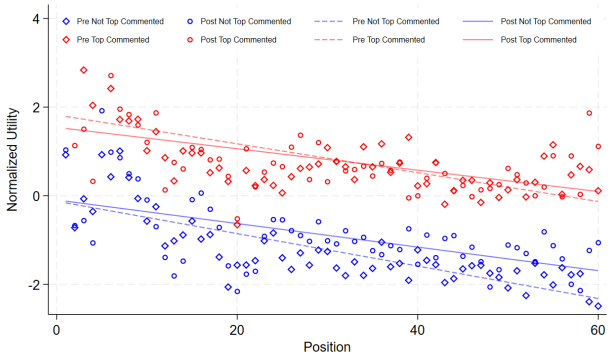
(D) Product Utility vs. Overall Pick



(E) Product Utility vs. Top Rating



(F) Product Utility vs. Top Commented



Notes: This figure summarizes key demand estimation results. Panel A presents the distribution of own-price elasticities across products and markets. Panel B compares average normalized product utility by position. Panel C to F compares average normalized product utility by position for Best Seller vs. Non Best Seller, Overall Pick vs. Non Overall Pick, Top Rating vs. Non Top Rating, and Top Commented vs. Non Top Commented.

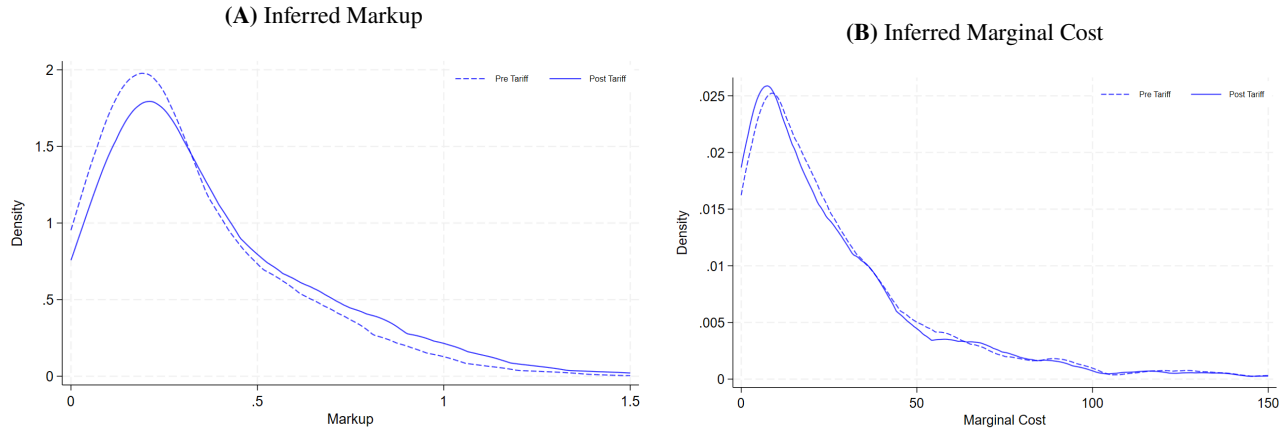
(the Lerner index) and the corresponding marginal cost.

6.2.1 Results

Figure 9 plots the distributions of implied markups and marginal costs before and after the tariff increase. Panel A shows a modest rightward shift in the markup distribution in the post-tariff period, consistent with sellers adjusting pricing margins after the tariff change. Panel B, by contrast, shows little visible change in the distribution of implied marginal costs.

These distributions should be interpreted as conditional on appearing in our observed sample of first-page search results. Products that exit the platform or no longer appear on the first page are not observed in the post-tariff period. The post-tariff densities therefore reflect equilibrium outcomes for products that remain active and visible, rather than the full set of products available before the tariff increase. This distinction is important because selection can mask underlying changes in seller costs.

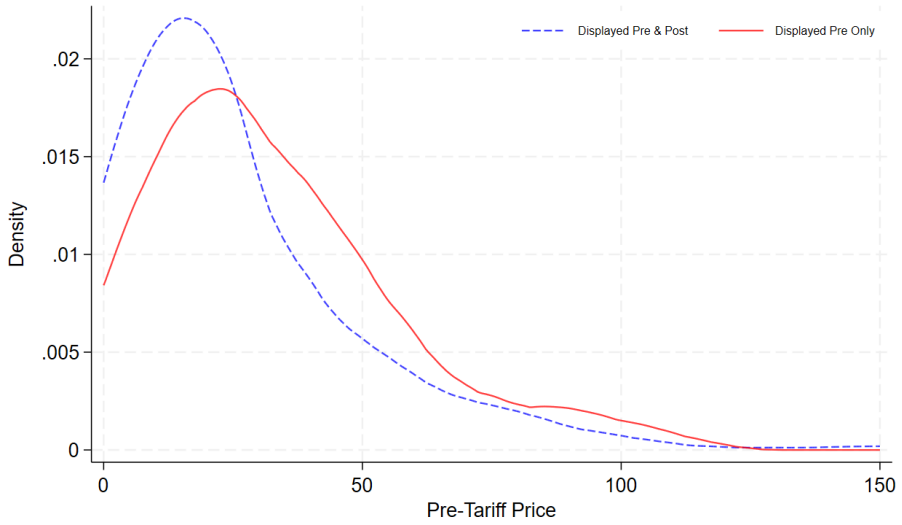
Figure 9. Seller Estimation Results



Notes: This figure summarizes key seller estimation results. Panel A presents the distribution of markup before and after the implementation of tariffs. Panel B presents the distribution of marginal costs before and after the implementation of tariffs.

In particular, products that disappear from the post-tariff sample are, on average, more expensive in the pre-tariff period than products that remain visible in both periods (see Figure 10). Assuming that higher pre-tariff prices are associated with higher marginal costs on average, the loss of these products from the observed sample truncates the upper part of the post-tariff marginal cost distribution. As a result, even if the tariff raises marginal costs at the product level, the observed cross-sectional distribution of implied marginal costs may show little movement because cost increases among surviving products are offset by the disappearance or demotion of relatively expensive products. For the same reason, we cannot conduct a clean pass-through analysis in this setting.

Figure 10. Pre-Tariff Price Distribution of High-Exposure Products



Notes: This figure compares the distribution of pre-tariff prices for high-exposure products that remain displayed in both the pre- and post-tariff periods and products that are displayed only in the pre-tariff period.

6.3 Overall Consumer Surplus Effects

We quantify the consumer surplus implications of the tariff shock and the contemporaneous ranking adjustment by computing consumer surplus implied by the estimated demand system. Under our logit framework, consumer surplus in market m during regime τ is given by

We quantify the consumer surplus implications of the tariff shock and the contemporaneous ranking adjustment by computing consumer surplus implied by the estimated demand system. Under our logit framework, consumer surplus in market m during regime τ is given by

$$CS_{m\tau} = \frac{1}{-\alpha_\tau} \ln\left(1 + \sum_{j \in m} \exp(\delta_{jm\tau})\right), \tag{6.7}$$

where α_τ is the period-specific price coefficient and $\delta_{jm\tau}$ denotes mean utility. This expression yields surplus per consumer in each keyword–day market. We compute surplus separately in the pre- and post-tariff periods using the corresponding estimated price sensitivity parameters.

To isolate the contribution of the ranking change, we construct a counterfactual that modifies only search exposure in the post period. Let $r_{jm\tau}$ denote realized position and $\hat{r}_{jm\tau}$ the position predicted under the pre-tariff

ranking rule. Using the estimated post-period position coefficient $\hat{\beta}_1^{pos}$, we define

$$\delta_{jm\tau}^{cf} = \begin{cases} \delta_{jm\tau}, & \text{Post} = 0, \\ \delta_{jm\tau} - \hat{\beta}_1^{pos} r_{jm\tau} + \hat{\beta}_1^{pos} \hat{r}_{jm\tau}, & \text{Post} = 1. \end{cases} \tag{6.8}$$

Counterfactual surplus is computed by replacing δ with δ^{cf} in Equation 6.7. The overall surplus change equals the difference between average post- and pre-period observed surplus. The ranking contribution is the gap between observed and counterfactual surplus in the post period, and the residual component captures all non-ranking adjustments. By construction, the total change equals the sum of these two components.

6.3.1 Results

Table 3 reports consumer surplus before and after the tariff, both under the observed ranking (i.e. status quo) and under the counterfactual in which the pre-tariff ranking rule is maintained.

Table 3. Consumer Surplus Effects

	Pre-Tariff Level ($) (1)	Post-Tariff Level ($) (2)	Change Level ($) (3)	Change Percent (%) (4)
Panel A: Overall				
Consumer Surplus (Status Quo)	0.511	0.527	0.016	3.06%
Consumer Surplus (Counterfactual)	0.511	0.586	0.075	14.66%
Panel B: High Exposure				
Consumer Surplus (Status Quo)	0.511	0.528	0.017	3.29%
Consumer Surplus (Counterfactual)	0.511	0.591	0.080	15.58%
Panel C: Low Exposure				
Consumer Surplus (Status Quo)	0.511	0.525	0.014	2.76%
Consumer Surplus (Counterfactual)	0.511	0.587	0.076	14.93%

Notes: This table compares outcomes under the status quo

Panel A of Table 3 shows that average consumer surplus rises from $0.511 in the pre-tariff period to $0.527 in the post-tariff period, an increase of $0.016, or 3.1 percent, under the observed ranking. Under the counterfactual that holds the ranking rule fixed, post-tariff consumer surplus would instead rise to $0.586, implying an increase of $0.075, or 14.7 percent, relative to the pre-tariff baseline. The gap between these two post-tariff outcomes, approximately $0.059, indicates that the contemporaneous ranking change offsets a substantial share of the consumer

surplus gains that would otherwise have followed the tariff shock.

Panels B and C show that this pattern is present in both high- and low-exposure categories. In high-exposure categories, consumer surplus increases by 3.3 percent under the observed ranking, compared with 15.6 percent under the counterfactual ranking. In low-exposure categories, the corresponding increases are 2.8 percent and 14.9 percent. Thus, in both sets of markets, consumer surplus would have risen considerably more in the absence of the ranking adjustment. The similarity of this pattern across exposure groups suggests that the welfare effect of the algorithmic change is not confined to the categories most directly exposed to Chinese products, but instead operates more broadly through the platform's reallocation of visibility across products.

Taken together, these results indicate that the tariff period is associated with an increase in measured consumer surplus, although this increase is substantially attenuated by the contemporaneous change in the platform's ranking algorithm. Two mechanisms may explain why consumer surplus rises despite the tariff shock. First, Amazon may have adjusted its marketplace in ways that helped retain consumers and preserve platform competitiveness during the tariff period. Second, products facing the largest cost increases may have exited or become less relevant, leaving a set of remaining products that was better able to remain competitively priced. However, the counterfactual exercise shows that consumer surplus would have been higher if the pre-tariff ranking rule had remained in place. This suggests that the post-tariff ranking adjustment reduced the consumer gains that would otherwise have emerged from the reallocation of products on the platform.

7 Conclusion

This paper studies how trade shocks propagate through digital marketplaces, focusing on the effects of heightened trade policy uncertainty and the sharp increase in U.S. tariffs on Chinese imports. Using high-frequency, product-level data from a major e-commerce platform, we show that trade policy affects not only prices, but also visibility, assortment, and consumer welfare. In particular, the adjustment to tariffs is shaped by the platform's ranking system, which determines how products are exposed to consumers and therefore how trade shocks are transmitted through the marketplace.

How do our results compare with previous findings in the literature? Prior studies (Amiti et al., 2019; Fajgelbaum et al., 2020; Cavallo et al., 2021) document near one-for-one pass-through of tariffs into duty-inclusive import prices, implying little offset in pre-tariff export prices by foreign sellers. By contrast, consistent with Nakamura and Zerom (2010) and Cavallo et al. (2025), we find incomplete pass-through in the context of digital retail markets. Our

results suggest that this more limited retail pass-through reflects the broader set of margins along which adjustment occurs on platforms, where cost shocks affect not only prices, but also product visibility, fulfillment, advertising, and competitive positioning.

Taken together, the empirical results show that trade shocks in digital markets propagate through multiple margins. Heightened trade policy uncertainty reduces the early-life visibility of new Chinese products, suggesting a worsening of the cold-start problem. Tariff increases generate only partial retail price pass-through, while also affecting Prime eligibility, badge allocation, sales rank, sponsored placement, and the composition of products shown to consumers. These changes imply that the effect of a trade shock on the platform cannot be understood solely through observed price adjustment; it also depends on how the platform reallocates attention across products.

The demand model sharpens this interpretation. We estimate that product position is strongly related to underlying demand, and that this relationship becomes steeper after the tariff shock. The same is true for platform-assigned badges and review volume, indicating that after the tariff, visibility becomes more tightly linked to demand-relevant product attributes. In the consumer surplus decomposition, we find that measured surplus rises in the post-tariff period, plausibly because the remaining products on the platform are those better able to remain competitive and because the platform has incentives to preserve consumer demand during the tariff episode. However, the counterfactual exercise shows that consumer surplus would have risen substantially more if the pre-tariff ranking rule had remained in place. Thus, the contemporaneous ranking adjustment attenuates the consumer gains that would otherwise have emerged.

More broadly, our findings highlight that in digital marketplaces, trade policy operates through an additional transmission channel: algorithmic allocation. Platforms do not merely transmit cost shocks mechanically into prices. They also determine which products consumers see, which products receive prominence, and how demand is redistributed across sellers. As a result, the incidence and welfare effects of trade shocks depend not only on pass-through, entry, and exit, but also on how platform algorithms map product characteristics into exposure. Understanding trade in the digital economy therefore requires integrating cost shocks with platform-mediated competition and algorithmic ranking.

References

Acemoglu, D. and Tahbaz-Salehi, A. (2025). The macroeconomics of supply chain disruptions. *Review of Economic Studies*, 92(2):656–695.

Amiti, M., Redding, S. J., and Weinstein, D. E. (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, 33(4):187–210.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Bown, C. P. (2025). US–China Trade War Tariffs: An Up-to-Date Chart. PIIE Charts, Peterson Institute for International Economics. Originally published September 20, 2019; updated May 14, 2025.

Brock, W. A. and Magee, S. P. (1978). The economics of special interest politics: The case of the tariff. *American Economic Review*, 68(2):246–250.

Cavallo, A., Gopinath, G., Neiman, B., and Tang, J. (2021). Tariff pass-through at the border and at the store: Evidence from us trade policy. *American Economic Review: Insights*, 3(1):19–34.

Cavallo, A., Llamas, P., and Vazquez, F. (2025). Tracking the short-run price impact of u.s. tariffs. Working Paper Active Draft, Harvard Business School. Version: August 7, 2025; Data through July 26, 2025.

De Loecker, J. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451.

Etro, F. (2023). Platform competition with free entry of sellers. *International Journal of Industrial Organization*, 89:102903.

Fajgelbaum, P., Goldberg, P., Kennedy, P., Khandelwal, A., and Taglioni, D. (2024). The us-china trade war and global reallocations. *American Economic Review: Insights*, 6(2):295–312.

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., and Khandelwal, A. K. (2020). The return to protectionism. *The Quarterly Journal of Economics*, 135(1):1–55.

Fajgelbaum, P. D. and Khandelwal, A. (2024). The value of de minimis imports. Technical report, National Bureau of Economic Research.

Fajgelbaum, P. D. and Khandelwal, A. K. (2022). The economic impacts of the us–china trade war. *Annual Review of Economics*, 14(1):205–228.

Fang, L., Chen, Y., Farronato, C., Yuan, Z., and Wang, Y. (2024). Platform information provision and consumer search: A field experiment.

Farronato, C., Fradkin, A., and MacKay, A. (2023). Self-preferencing at amazon: evidence from search rankings. In *AEA Papers and Proceedings*, volume 113, pages 239–243. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Flaaen, A., Hortaçsu, A., and Tintelnot, F. (2020). The production relocation and price effects of us trade policy: the case of washing machines. *American Economic Review*, 110(7):2103–2127.

Fleck, A. (2024). "made in china, sold on amazon – share of items sold on amazon by country of origin". Statista Chart. Accessed 2025-11-06.

FTC (2023). Ftc complaint against amazon. Technical Report Case No.: 2:23-cv-01495, Federal Trade Commission.

General Administration of Customs of the People's Republic of China (2026). China customs statistical data query platform. http://stats.customs.gov.cn/. Online database. Accessed January 13, 2026.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.

Grossman, G. M., Helpman, E., and Redding, S. J. (2024). When tariffs disrupt global supply chains. *American Economic Review*, 114(4):988–1029.

Gutierrez, G. (2021). The welfare consequences of regulating amazon. *Available at SSRN 3965566*.

He, S. and Hollenbeck, B. (2020). Sales and rank on amazon. com. *Available at SSRN 3728281*.

Hiller, R. S. and Savage, S. J. (2021). Tariff pass-through and welfare in the tablet computer market. *The Journal of Industrial Economics*, 69(2):369–409.

Leith, J. C. (1971). The effects of tariffs on production, consumption, and trade: A revised analysis. *American Economic Review*, 61(1):74–81.

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

Nakamura, E. and Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, 77(3):1192–1230.

Nassauer, S. and Khan, N. (2025). The biggest retailers are thriving in the tariff economy. *The Wall Street Journal*. Accessed via The Wall Street Journal on September 17, 2025.

Teh, T.-H. (2022). Platform governance. *American Economic Journal: Microeconomics*, 14(3):213–254.

Yu, C. (2024). The welfare effects of sponsored product advertising. *Available at SSRN 4817542*.

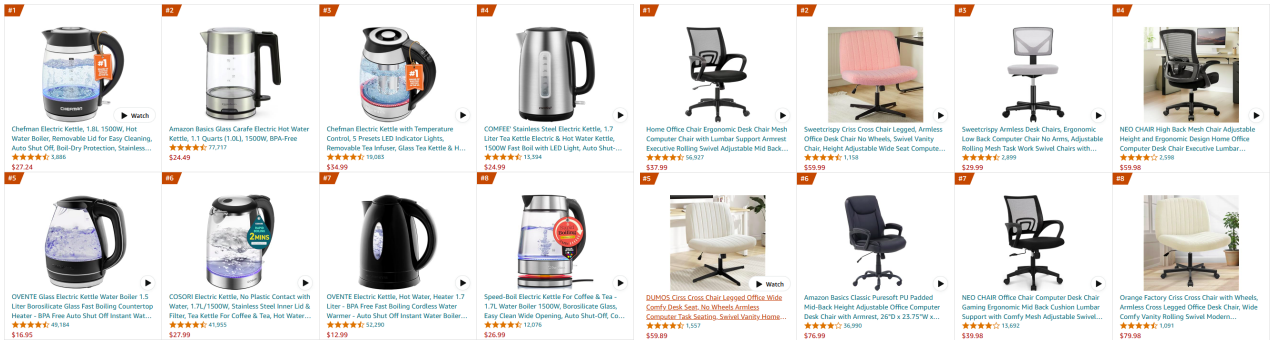
Appendices

A.1 Product Information

We illustrates examples of top-selling products within narrow subcategories of the "Home & Kitchen" product category on Amazon.

Figure A.1. Example of Product Similarity within Subcategory

(A) Example 1 **(B)** Example 2



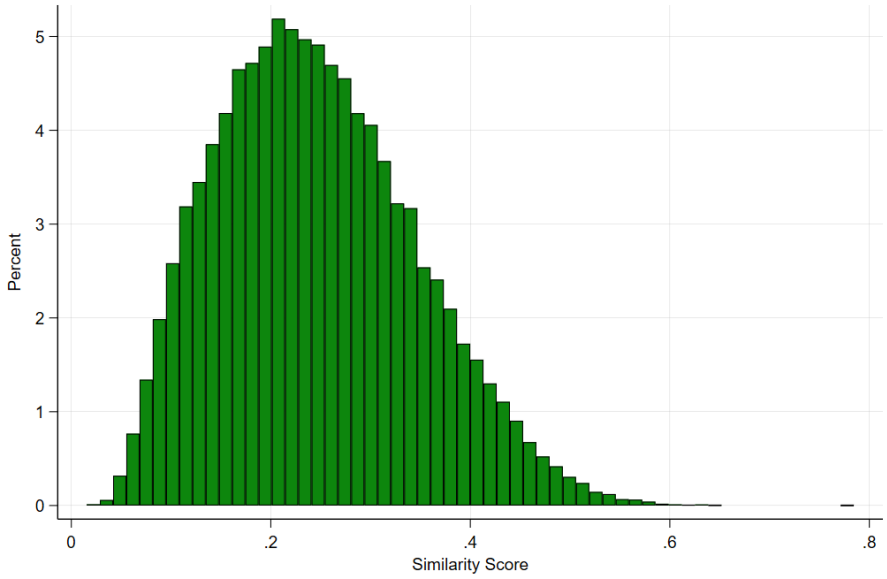
Notes: The figure illustrates examples of top-selling products within narrow subcategories of the "Home & Kitchen" product category on Amazon. Panel A shows products from the "Electric Kettle" subcategory, while panel B displays products from the "Office Chair" subcategory.

The algorithm we implement tokenizes titles into individual words and co-located word pairs. Then, it computes a similarity score between zero and one by comparing each product's token set with those of all other products on the same search results page for a given keyword. We present the distribution of similarity scores for products listed on the first search results page. Figure A.2 illustrates the degree of similarity between products based on their descriptions, showing how closely they align with others displayed for the same keyword on the first results page.

Figure A.3 provides an example of products with low and high similarity scores for a specific keyword. The item in the green box has a similarity score of about 0.7 while the item in the red ox has a similarity score of about 0.1.

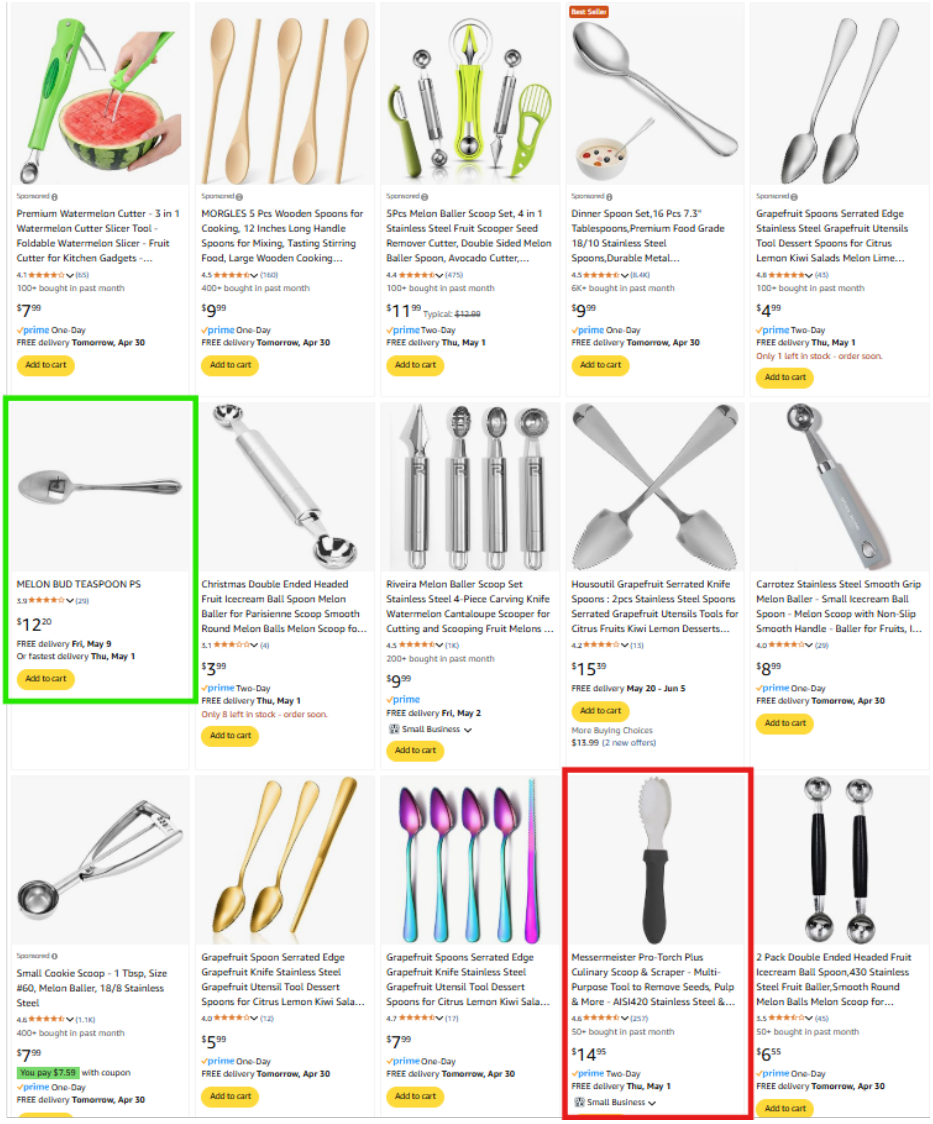
Figure A.4 provides examples of the interfaces from which we scrape product information.

Figure A.2. Similarity Score Distribution



Notes: This figure illustrates the distribution of the similarity scores of products on the first search results page.

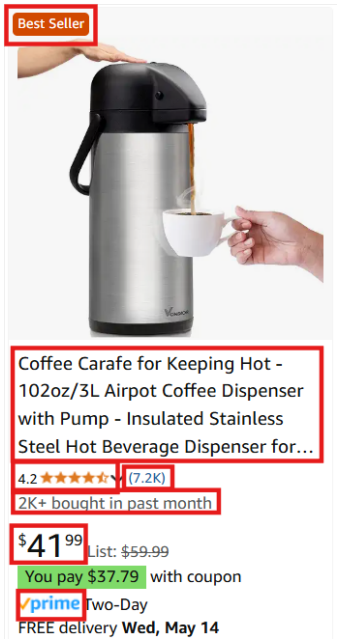
Figure A.3. Similarity Score Example



Notes: This figure illustrates an example of a product with a low similarity score (red box) and a high similarity score (green box) with the first result page of the keyword "Melon Spoons".

Figure A.4. Example of Product Information

(A) Primary Information



(B) Secondary Information

Brand	VONDIOR
Material	Stainless Steel
Color	1 Pack
Capacity	3 Liters
Special Feature	Airpot
Product Dimensions	15.6"W x 5.6"H
Included Components	Airpot
Shape	Airpot
Item Weight	1.6 Kilograms
With Lid	Yes
UPC	709202435534 680491776051
Item Weight	3.52 pounds
Manufacturer	VONDIOR
ASIN	B01MRGIZE9
Country of Origin	China
Item model number	0
Customer Reviews	4.2 (7,284) 4.2 out of 5 stars
Best Sellers Rank	#7,857 in Home & Kitchen (See Top 100 in Home & Kitchen) #1 in Thermal Carafes
Is Discontinued By Manufacturer	No
Date First Available	November 6, 2016

Notes: The figure provides examples of product information extracted from online listings, distinguishing between primary attributes (e.g., title) and secondary attributes (e.g., seller details).

A.2 Summary Statistics

Table A.1. Summary Statistics

	Obs. (1)	Min. (2)	Mean (3)	Median (4)	Max. (5)	S.D. (6)
Sales Rank	1,195,986	1.00	160,996.00	35,183.00	13,118,728.00	431,543.80
Sales Last Month	880,732	50.00	1,362.93	400.00	100,000.00	3,918.82
Price (US $)	1,209,818	0.00	49.81	23.99	999.99	85.64
Referral Fee (%)	1,180,495	4.93	14.81	15.01	45.03	1.14
Star Rating (1-5)	1,168,749	1.00	4.46	4.50	5.00	0.33
Comments	1,080,430	0.00	3,370.63	373.00	849,510.00	13,529.69
Prime	1,226,532	0.00	0.73	1.00	1.00	0.44

Notes: This table reports summary statistics. A lower sales rank corresponds to higher sales. Sales last month is a categorical variable indicating the product's sales volume in the previous month. Price is measured in U.S. dollars. Referral fee is expressed as a percentage. Star Rating is an ordinal measure ranging from 1 (lowest) to 5 (highest). Comments refers to the number of user-generated product reviews. Prime is a binary indicator equal to 1 if the product is Amazon Prime and 0 otherwise

A.3 Personalization

Figure A.5. Personalization Example



Notes: This figure illustrates variation in product rankings for the same search keyword across different user sessions. Highlighted boxes indicate identical products appearing in different positions.

A.4 Product Distribution by Country

Table A.2. Distribution by Country of Origin

Country	No.	%	Country	No.	%
Afghanistan	95	0.022%	Korea, Republic of	680	0.159%
Albania	47	0.011%	Kyrgyzstan	19	0.004%
Algeria	8	0.002%	Lao People's D.R.	22	0.005%
American Samoa	24	0.006%	Latvia	15	0.004%
Andorra	29	0.007%	Lithuania	5	0.001%
Anguilla	72	0.017%	Macau	14	0.003%
Antarctica	7	0.002%	Malaysia	677	0.158%
Argentina	36	0.008%	Mexico	1380	0.322%
Armenia	55	0.013%	Morocco	225	0.053%
Australia	34	0.008%	Namibia	1	0.000%
Austria	106	0.025%	Netherlands	154	0.036%
Bangladesh	56	0.013%	New Zealand	27	0.006%
Belgium	67	0.016%	Niue	9	0.002%
Bouvet Island	2	0.000%	Pakistan	3359	0.784%
Brazil	213	0.050%	Philippines	18	0.004%
Brunei Darussalam	2	0.000%	Poland	442	0.103%
Bulgaria	42	0.010%	Portugal	175	0.041%
Cambodia	61	0.014%	Puerto Rico	20	0.005%
Canada	453	0.106%	Romania	355	0.083%
C.A.R.	25	0.006%	Russian Federation	13	0.003%
Chile	147	0.034%	Singapore	139	0.032%
China	370349	86.438%	Slovakia	172	0.040%
Christmas Island	4	0.001%	Slovenia	83	0.019%
Colombia	256	0.060%	South Africa	1	0.000%
Croatia	6	0.001%	Spain	424	0.099%
Czech Republic	74	0.017%	Sri Lanka	34	0.008%
Denmark	76	0.018%	St. Helena	4	0.001%
Dominican Republic	1	0.000%	Sweden	107	0.025%
Ecuador	4	0.001%	Switzerland	624	0.146%
Egypt	164	0.038%	Taiwan	3054	0.713%
El Salvador	47	0.011%	Thailand	1595	0.372%
Estonia	17	0.004%	Tunisia	165	0.039%
F.S. of Micronesia	13	0.003%	Turkey	1624	0.379%
Finland	17	0.004%	USA	22617	5.279%
France	691	0.161%	Ukraine	522	0.122%
Georgia	18	0.004%	United Arab Emirates	38	0.009%
Germany	1129	0.264%	United Kingdom	160	0.037%
Ghana	5	0.001%	Venezuela	24	0.006%
Greece	7	0.002%	Vietnam	3654	0.853%
Hong Kong	234	0.055%	Hungary	106	0.025%
India	7286	1.701%	Ireland	4	0.001%
Indonesia	617	0.144%	Israel	373	0.087%
Italy	1273	0.297%	Jamaica	19	0.004%
Japan	1432	0.334%			

Notes: This table shows the number and percentage of products by country of origin.

A.5 Timeline of Major U.S. Tariff Measures in 2025

We present the major tariffs changes in 2025 between China and the U.S. in Table A.3.

Table A.3. Major Trade Policy Events Between the U.S. and China in 2025

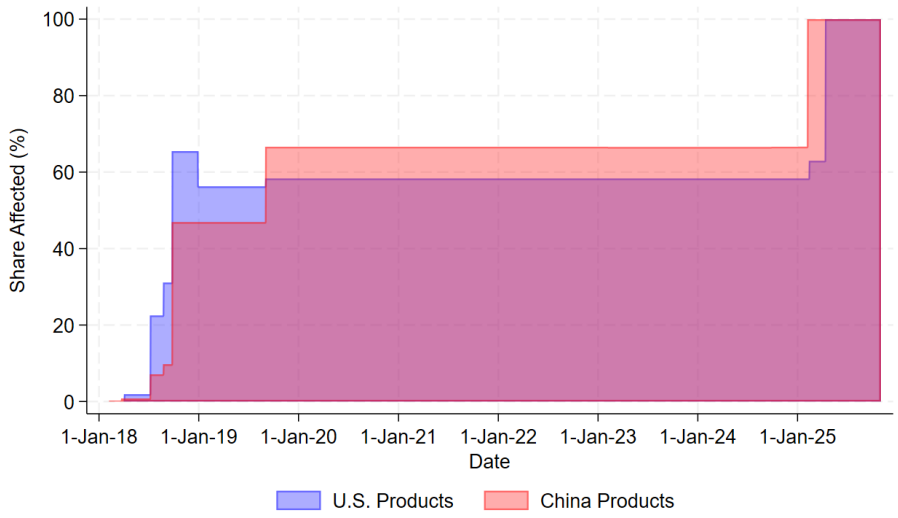
Date	Country	Event
February 4, 2025	US	10% tariff imposed on all imports from China.
February 10, 2025	China	China imposes retaliatory tariffs.
March 4, 2025	US	Tariffs on imports from China increased from 10% to 20%.
March 10, 2025	China	China responds with additional tariffs.
April 10, 2025	US	Tariffs on imports from China raised to 125%.
April 10, 2025	China	China enacts corresponding tariff increases.
May 14, 2025	US	Tariffs on imports from China reduced from 125% to 10%.
May 14, 2025	China	Tariffs on U.S. imports reduced from 125% to 10%.

Notes: This table summarizes key 2025 trade policy actions between the U.S. and China. Source: Trump's Trade War Timeline 2.0: An Up-to-Date Guide.

A.6 Share of Affected Products

Figure A.6 shows the share of Chinese products subject to retaliatory tariffs. By February 2025, nearly all Chinese exports to the U.S. were covered by tariffs.

Figure A.6. Affected Share Trends Timeline

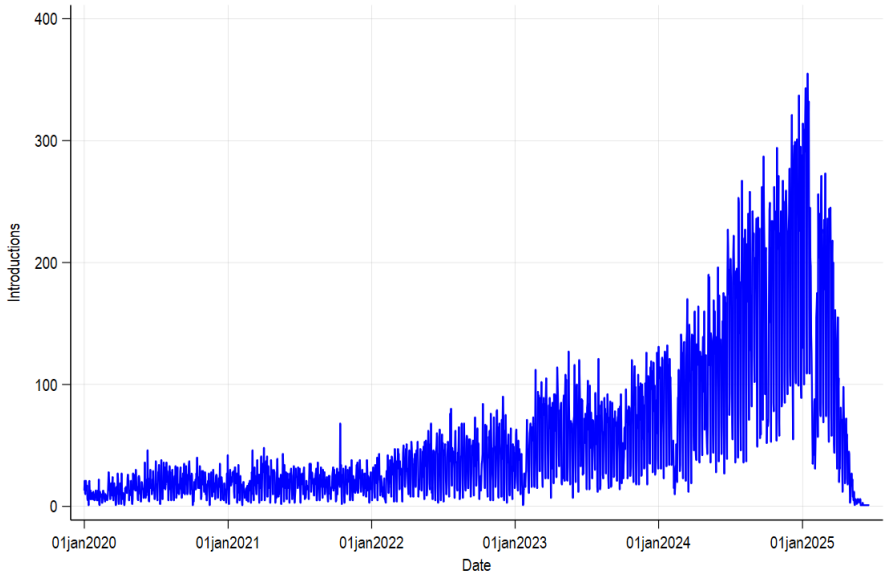


Notes: This figure plots the share of Chinese affected products by U.S. tariffs and the the share of U.S. affected products by Chinese tariffs over time. Data from Bown (2025)

A.7 Available Since Time Series

Figure A.7 presents the time series of the number of products available since on each date on the first search results page.

Figure A.7. Available Since Time Series



Notes: This figure plots the number of products available since a specific date.

A.8 Robustness Check: Pre-Tariff Shock Regression

To assess the validity of the parallel trends assumption underlying our DiD, we estimate a pre-tariff shock regression of variables of interest on the interaction between the China indicator and the time dummies, including keyword fixed effects:

$$y_{ct} = \alpha + \beta \textit{High-Exposure}_c + \gamma \textit{Generation}_t + \delta(\textit{High-Exposure}_c \times \textit{Generation}_t) + \kappa_c + \varepsilon_{ct}. \quad \text{(A.1)}$$

This specification isolates differences between High-Exposure and Low-Exposure product categories within each generation prior to the tariff shock. The pre-tariff shock regression results in Table A.4 indicate that, across generations, there are generally no statistically significant differences in prices or most other product characteristics between China and non-China products prior to the tariff, supporting the plausibility of the parallel trends assumption.

Table A.4. Pre-Tariff Shock Regression Results

	Price	Prime Delivery	Best Seller	Overall Pick	Sales Rank	Ratings	Comments	Brand Variety	Manuf. Variety	# Sponsored
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High-Exposure	0.013	0.005	-0.004	-0.005**	0.046	0.000	-0.011	0.110***	0.104***	0.195***
	(0.014)	(0.004)	(0.003)	(0.002)	(0.038)	(0.001)	(0.029)	(0.026)	(0.026)	(0.039)
High-Exposure × *t=2*	-0.013	0.003	0.001	0.002*	-0.045	0.000	0.021	-0.027*	-0.026*	-0.076***
	(0.011)	(0.003)	(0.002)	(0.001)	(0.033)	(0.001)	(0.023)	(0.015)	(0.016)	(0.027)
High-Exposure × *t=3*	0.005	0.002	-0.000	0.003**	-0.004	0.000	0.003	-0.041***	-0.041***	-0.080***
	(0.010)	(0.004)	(0.002)	(0.001)	(0.035)	(0.001)	(0.027)	(0.013)	(0.013)	(0.028)
High-Exposure × *t=4*	-0.016	-0.002	0.002	0.003**	0.051	0.000	0.008	-0.046***	-0.043***	-0.097***
	(0.011)	(0.004)	(0.002)	(0.001)	(0.035)	(0.001)	(0.027)	(0.014)	(0.014)	(0.028)
High-Exposure × *t=5*	-0.025**	0.002	-0.000	0.002	0.016	0.001	0.013	-0.035***	-0.036***	-0.059*
	(0.012)	(0.004)	(0.002)	(0.001)	(0.035)	(0.001)	(0.028)	(0.013)	(0.013)	(0.031)
High-Exposure × *t=6*	-0.010	-0.001	0.000	0.002**	-0.045	-0.000	0.021	-0.052***	-0.046***	-0.092***
	(0.012)	(0.004)	(0.002)	(0.001)	(0.035)	(0.001)	(0.028)	(0.013)	(0.013)	(0.029)
High-Exposure × *t=7*	-0.004	-0.001	-0.001	0.002*	0.014	0.000	0.019	-0.032**	-0.030**	-0.090***
	(0.014)	(0.004)	(0.002)	(0.001)	(0.040)	(0.001)	(0.027)	(0.014)	(0.014)	(0.028)
High-Exposure × *t=8*	-0.012	0.003	0.001	0.002	-0.043	-0.000	0.057*	-0.040***	-0.041***	-0.102***
	(0.013)	(0.004)	(0.002)	(0.001)	(0.042)	(0.001)	(0.030)	(0.014)	(0.014)	(0.033)
High-Exposure × *t=9*	0.002	-0.003	-0.001	0.003***	-0.035	-0.000	0.003	-0.045***	-0.041***	-0.133***
	(0.013)	(0.004)	(0.002)	(0.001)	(0.038)	(0.001)	(0.030)	(0.014)	(0.014)	(0.031)
High-Exposure × *t=10*	-0.004	-0.006	0.001	0.001	0.023	0.000	0.051	-0.035**	-0.031**	-0.109***
	(0.012)	(0.004)	(0.002)	(0.001)	(0.041)	(0.001)	(0.036)	(0.014)	(0.014)	(0.029)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12820	12820	12820	12820	12820	12820	12817	12820	12820	12820

Notes: This table shows the estimates from Equation A.1. γ estimates are excluded in the table. Robust standard errors are in parenthesis. * p<0.1, ** p<0.05, *** p<0.01.

A.9 Robustness Check: Product Level Analysis

We conduct a robustness check by replicating our estimation at the product level. This allows us to verify that our findings are not driven by alternative channels and remain consistent across level of granularity.

We re-estimate Equation 5.2 at the product level.

Table A.5. Tariff Transmission Results

	Price	Prime Delivery	Best Seller	Overall Pick	Sales Rank	Rating (1–5)	Com-ments
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post	0.009***	-0.339***	0.003	-0.001	0.053***	-0.001***	-0.511***
	(0.001)	(0.003)	(0.002)	(0.001)	(0.003)	(0.000)	(0.011)
High-Exposure × *Post*	0.007***	-0.035***	-0.003	0.001	0.022***	-0.000	0.180***
	(0.001)	(0.003)	(0.002)	(0.001)	(0.003)	(0.000)	(0.012)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ASIN FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	609288	615302	615302	615302	604447	587990	549543

Notes: This table shows the estimates from Equation 5.2 on a product-level. Standard errors are clustered at the ASIN-level and in parenthesis. * p<0.1, ** p<0.05, *** p<0.01.

A.10 Robustness Check: *De Minimis* Termination

As a robustness check, we exploit the termination of the *de minimis* exemption, which directly affected products valued under $800 from China and Hong Kong that were previously allowed to enter the U.S. duty-free. With the exemption removed, these products became subject to a tariff of either 120% of their declared value or a flat fee of $100 per item, whichever is higher.[11] This sudden and substantial change serves as an exogenous tariff shock to identify the impact of trade policy.

We re-estimate Equation 5.2 using May 2, 2025—the date the *de minimis* exemption was revoked—as the timing of the exogenous shock.

Table A.6. Tariff Results

	Price	Prime Delivery	Best Seller	Overall Pick	Sales Rank	Ratings	Com-ments	Brand Variety	Manuf. Variety	# Sponsored
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
High-Exposure	-0.001	0.052***	-0.007***	-0.003***	0.103***	-0.001	-0.031	0.169***	0.164***	0.198***
	(0.009)	(0.009)	(0.002)	(0.001)	(0.027)	(0.001)	(0.043)	(0.022)	(0.021)	(0.026)
Post	-0.045***	-0.500***	0.003***	-0.000	-0.041**	0.000	0.240***	-0.013	-0.020*	-0.005
	(0.007)	(0.007)	(0.001)	(0.000)	(0.018)	(0.000)	(0.017)	(0.011)	(0.011)	(0.012)
High-Exposure × *Post*	0.017**	-0.035***	-0.003***	-0.000	0.043**	-0.001	-0.014	0.033***	0.037***	0.019
	(0.008)	(0.008)	(0.001)	(0.000)	(0.020)	(0.001)	(0.019)	(0.011)	(0.011)	(0.012)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782	21,782

Notes: This table shows the estimates from Equation 5.2. Columns (1) shows results related to price, Columns (2)-(10) show results related to delivery, platform assigned badges, sales, ratings, comments, brand variety, manufacturer variety, and number of sponsored products. Robust standard errors are clustered at the ASIN-level and in parenthesis. Statistical significance is denoted by * p<0.1, ** p<0.05, and *** p<0.01.

[11] After May 14, 2025, these rates were reduced to 54% and $100, respectively.

A.10.1 Product Level Analysis

We conduct an additional robustness check by replicating our analysis using the termination of the *de minimis* exemption. We re-estimate Equation 5.2 using May 2, 2025—the date the *de minimis* exemption was revoked—as the timing of the exogenous shock and at the product level.

Table A.7. Tariff Transmission Results

	Price	Prime Delivery	Best Seller	Overall Pick	Sales Rank	Rating (1–5)	Comments
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post	0.006***	-0.534***	0.001	-0.001	0.062***	-0.001***	1.071***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.000)	(0.011)
High-Exposure × *Post*	0.012***	-0.016***	-0.000	0.000	0.034***	-0.000***	-1.050***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.000)	(0.012)
Keyword FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ASIN FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	609288	615302	615302	615302	604447	587990	549543

Notes: This table shows the estimates from Equation 5.2 on a product-level. Standard errors are clustered at the ASIN-level and in parenthesis. * p<0.1, ** p<0.05, *** p<0.01.

A.11 Demand Estimation Details

A.11.1 Calibration of Fraction of Sales From Focal Keywords

The demand model derives each product's market share within a specific keyword. However, the data includes platform-wide sales, and products often appear in search results for multiple keywords. Ignoring this can bias demand estimates. For example, a product that ranks poorly in the observed keyword may be excluded from many consumers' consideration sets but still generate substantial sales by ranking higher under other, more relevant keywords. Attributing all its sales to the observed keyword would overstate its utility.

To address this issue, we calibrate the fraction of a product's total sales attributable to the observed keyword in our dataset, denoted $\kappa \in [0, 1]$. Specifically, we record the search volume of each keyword in which a product appears and assume that the ratio of observed keyword search volume to the product's total search volume translates into the corresponding fraction of the product's total sales generated through that keyword. Let q_{it} denote products j's total sales on day t, and M_t the market size of the observed keyword. We compute the market share for each product as follows:

$$s_{jt}(p) = \frac{\kappa_j q_{jt}}{M_t} \tag{A.1}$$

A.11.2 Calibration of Market Size

In this section, we describe the calibration of the daily keyword-specific market size M_t. Amazon accounts for about 40% of U.S. online purchases which does not necessary translate to the search level. We calibrate the market size using aggregate data on *(i)* total units sold on Amazon and *(ii)* the total number of searches on Amazon.

First, Amazon's Q1 and Q2 2025 earnings release indicates that online stores generated on average $59.4 billion revenues, with 60% attributed to North America. Our analysis focuses on Amazon.com, a domain primalrly serving the U.S. market. Using sales weighted avergage price of $49.81 from my sample, we estimate 0.72 billion units of first-party products were sold in North America on average in a quarter. Since 61.5% of the units were sold by third-party seller, total marketplace sales are approximately $\frac{0.72}{61.5\%} = 1.2$ billion units.

Second, Amazon.com receives roughly 2.3 billion visits on average in the first quarter in 2024 (we assume the same for Q1 and Q2 in 2025). Assuming an average of two keywords searches per visit, total quarterly keyword searches are estimated at 13.8 billion. Combining these statistics yields an average purchase incidence (i.e. inside

market share) of 8.3% (1.2 billion units sold divided by 13.8 billion searches). Accordingly, we calibrate the market size for each keyword on each day as:

$$M_t = \frac{1}{0.083} \sum_{j \in J} \kappa_{jt} q_{it} \tag{A.2}$$

A.11.3 First-Stage Results

Table A.8. First Stage Estimates

	Price
	(1)
Post=0 × *IV*	0.00056***
	(0.00009)
Post=1 × *IV*	0.00027***
	(0.00009)
F-Statistics	28.74 / 28.09
Controls	Yes
ASIN FE	Yes
Keyword FE	Yes
Observations	16,046

Notes: This table reports the first stage results. Standard errors are clustered at the ASIN level and reported in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table A.9. Demand Parameter Estimates

	(1) OLS	(2) IV
Price × Pre	-0.033***	-0.169***
	(0.007)	(0.042)
Price × Post	-0.033***	-0.143***
	(0.006)	(0.041)
Best Seller × Pre	0.801***	0.688***
	(0.116)	(0.129)
Best Seller × Post	0.725***	0.812***
	(0.167)	(0.171)
Overall Pick × Pre	0.663***	0.531***
	(0.179)	(0.170)
Overall Pick × Post	0.810***	0.869***
	(0.134)	(0.146)
Rating × Pre	0.462	1.106*
	(0.434)	(0.571)
Rating × Post	0.535	1.006*
	(0.434)	(0.572)
Comments × Pre	0.000**	0.000**
	(0.000)	(0.000)
Comments × Post	0.000*	0.000*
	(0.000)	(0.000)
Position × Pre	-0.012***	-0.008***
	(0.002)	(0.002)
Position × Post	-0.011***	-0.010***
	(0.002)	(0.003)
Observations	16,377	16,046

Notes: This table reports the demand parameters estimates. Standard errors are clustered at the ASIN level and reported in parentheses. * p<0.1, ** p<0.05, *** p<0.01.