# Private Labels and Product Imitation: Evidence from Amazon*

Malika Korganbekova
*University of Chicago Booth School of Business*

Alibek Korganbekov
*Boston University*

March 26, 2026

**Abstract**

Digital platforms that operate as both marketplace and seller can leverage proprietary seller data to introduce competing products, potentially harming the third-party sellers they host. Proposed regulations—including the EU Digital Markets Act and U.S. platform competition legislation—seek to restrict this practice, yet systematic evidence on its prevalence and effects is scarce. We study Amazon's private-label program by scraping over 770 million product pages across Amazon's four largest marketplaces (U.S., Germany, U.K., and Canada) and developing image recognition methods to measure product imitation at scale. We find that 73% of Amazon's private-label products closely resemble pre-existing third-party products, with imitation disproportionately targeting commercially successful products sold by established sellers. For consumers, private-label products are 41% cheaper on average and carry higher ratings. For third-party sellers, the effects are negative: targeted products experience worsening sales ranks and higher discontinuation rates, while at the category level, private-label entry significantly reduces the number of products, drives down the highest prices, and erodes the cumulative review stock—without degrading product quality as measured by star ratings. The category-level effects are consistent across all four marketplaces; the imitation and product-level analyses are conducted using U.S. data.

**JEL Codes:** L13, L41, L81, D43, O31
**Keywords:** Private labels, digital platforms, product imitation, marketplace competition, Amazon, third-party sellers

*We thank conference participants at the Kellogg School of Management, Yale University AI Center, Marketing Science 2025 for their useful feedback and comments. Correspondence: malika.korganbekova@chicagobooth.edu, alibekk@bu.edu.

# 1 Introduction

A growing number of digital platforms operate in dual mode: they simultaneously serve as a marketplace connecting consumers and third-party sellers, and as a seller of their own products on that same marketplace. Amazon, Apple, Walmart, and Target all introduce private-label products that compete directly with the third-party sellers who depend on their platforms. This dual role has attracted intense regulatory scrutiny, most notably in the U.S. House Committee on the Judiciary's 2020 Investigation of Competition in Digital Markets, which devoted over one hundred pages to the concern that platforms exploit proprietary marketplace data to introduce competing products—a practice commonly referred to as *sherlocking* (U.S. House Committee on the Judiciary, 2020). Despite the prominence of this debate, systematic empirical evidence on the extent of product imitation by platforms, the mechanisms underlying it, and its consequences for consumers and third-party sellers remains scarce. Existing empirical work relies on case studies or limited product samples; measuring imitation at scale requires automated similarity methods that can compare hundreds of millions of products, which we develop in this paper.

This paper addresses two central questions. First, is there systematic evidence that Amazon imitates third-party seller products when introducing its private labels, and what does the imitation pattern reveal about the platform's entry strategy? Second, what is the impact of Amazon's private-label entry on consumer welfare and on the outcomes of affected third-party sellers?

The welfare implications of private-label entry by a dominant platform are theoretically ambiguous. Consumers may benefit from similar products at lower prices (Farronato et al., 2025), but private-label entry may reduce third-party sellers' revenues and weaken their incentives to invest in quality and innovation. If the threat of imitation deters entry or reduces quality, the long-run consequences for consumers could be negative even if short-run price effects are favorable. While we do not estimate

total welfare, our reduced-form estimates provide evidence on key components of this tradeoff—prices, product variety, and seller survival—that are essential for informing the ongoing policy debate.

To answer these questions, we assemble a novel dataset of unprecedented scale. We scrape over 770 million product pages from Amazon's four largest marketplaces—the U.S., Germany, the U.K., and Canada—covering both private-label and third-party seller products from 2018 to the present. For each product, we observe prices, ratings, review counts and distributions, product images, textual descriptions, sales ranks, introduction dates, and product taxonomy information. This panel structure allows us to track product entry, measure similarity between products, and estimate the causal effects of private-label entry on incumbent sellers over time. We identify causal effects using a combination of similarity-based matching for product-level estimates and staggered difference-in-differences designs that exploit variation in the timing of private-label entry across categories for category-level estimates.

We develop a methodology to systematically measure product similarity at scale. Our primary approach uses visual similarity: we train a Siamese neural network that maps product images into a 512-dimensional embedding space, where cosine similarity captures visual resemblance. We complement this with textual similarity based on sentence embeddings of product titles and descriptions; textual results corroborate the visual findings and are available upon request. These tools allow us to identify, for each of the 1,673 Amazon private-label products in our 14 focal product categories, which third-party products—out of more than 6 million candidates—are most similar.

Our analysis yields four main findings. First, we document substantial visual similarity between Amazon's private-label products and pre-existing third-party seller products. Out of 1,673 private-label products, 1,222 (73%) have at least one third-party predecessor with a visual cosine similarity of 0.9 or above—a threshold corresponding to near-identical appearance. The products Amazon imitates are cheaper, have significantly

more reviews (particularly five-star reviews), hold better sales ranks, and are sold by larger, more established sellers. While high visual similarity is consistent with deliberate imitation, it may also reflect convergent design in categories where functional constraints limit the space of viable products. Our analysis therefore focuses on systematic selection patterns—which products Amazon targets and how this targeting differs from third-party imitation—rather than treating similarity as proof of intentional copying. The selection patterns we document are consistent with, but not definitive evidence of, a data-informed entry strategy; we cannot rule out that Amazon simply targets publicly observable signals of commercial success more effectively than other market participants.

Second, Amazon's imitation behavior differs qualitatively from imitation among third-party sellers. Third-party imitators target products with lower prices and do not systematically select on sales rank or seller characteristics to the same degree. This asymmetry admits two interpretations: Amazon's more targeted selection could reflect access to proprietary data unavailable to ordinary sellers, or third-party sellers may simply pursue different competitive strategies when entering with similar products.

Third, Amazon's private-label products are, on average, $12 (41%) cheaper than their most similar third-party counterparts, carry higher ratings, and accumulate substantially more reviews—despite entering the market an average of 774 days later. These patterns suggest that private-label entry expands the consumer choice set and exerts downward pressure on prices.

Fourth, we estimate the causal impact of private-label entry on third-party sellers at both the product and category levels. At the product level, third-party products directly targeted by Amazon's imitation experience a statistically significant deterioration in sales rank (4.5%), downward price adjustments (1.1%), and an elevated probability of discontinuation. At the category level, using the Callaway and Sant'Anna (2021) difference-in-differences estimator, which is robust to heterogeneous treatment effects in staggered adoption designs, we find that private-label entry is associated with a

significant relative decline in the number of products in treated categories (treated categories have 37–57% fewer products than they would have absent Amazon's entry, when restricting to categories with at least 50 baseline products), drives down the highest prices in treated categories (maximum price declines by approximately 39–61%), reduces listing investment, and erodes the cumulative review stock—while leaving product quality as measured by star ratings largely unaffected. These effects are consistent across all four Amazon marketplaces we study (U.S., Germany, U.K., Canada).

This paper contributes to several strands of the literature. A large theoretical literature examines platform self-preferencing and private-label introduction (Farrell and Katz, 2000; Hagiu and Spulber, 2013; Hagiu et al., 2015, 2022; Jiang et al., 2011; Anderson and Bedre-Defolie, 2025; Choi et al., 2025; Hervas-Drane and Shelegia, 2022; Madsen and Vellodi, 2022; Zou and Zhou, 2023). Our paper provides large-scale empirical evidence on several mechanisms these models highlight, including the role of data in guiding entry decisions and the heterogeneous effects on incumbent sellers.

On the empirical side, our work is most closely related to studies of Amazon's private-label strategy (Chen and Tsai, 2022; Farronato et al., 2022; Gutierrez, 2021; Jeffries and Yin, 2021; Lam, 2024; Lee and Musolff, 2021; Raval, 2022; Crawford et al., 2022). We make two primary contributions. First, we provide the first systematic, large-scale analysis of product imitation using image recognition methods applied to over 770 million products across four marketplaces, moving beyond case studies to document the prevalence and characteristics of imitation across a broad set of categories. Second, we provide direct evidence on competitive effects at both the category and product levels—including seller exit, price compression, listing quality, review accumulation, and individual product discontinuation—using heterogeneity-robust estimators applied across four independent Amazon marketplaces, providing the most comprehensive causal evidence to date on how private-label entry reshapes platform competition.

The remainder of the paper is organized as follows. Section 2 describes the empirical

setting. Section 3 presents our data and similarity measures. Section 4 documents the evidence on product imitation. Section 5 examines consumer-side outcomes. Section 6 estimates product-level effects on targeted third-party sellers. Section 7 estimates category-level effects of private-label entry. Section 8 concludes.

## 2 Empirical Setting

Private-label products—goods manufactured by or on behalf of a retailer and sold under the retailer's own brand—have been a feature of retail since the 1950s. Grocery chains pioneered the practice, and general merchandise retailers followed: Target's "Up & Up," Costco's "Kirkland Signature," and Walmart's "Great Value" all compete directly with branded products on their shelves. By eliminating brand marketing costs and leveraging shelf-space control, retailers offer comparable products at lower prices while earning higher margins (Dubé, 2024).

Digital marketplaces introduce a fundamentally different structure. On platforms such as Amazon, third-party sellers list products, pay referral and fulfillment fees, and compete for visibility through search and recommendation algorithms. The platform observes granular data on every dimension of seller and consumer behavior: search queries, click-through rates, conversion rates, return rates, pricing histories, and inventory levels. This informational advantage is orders of magnitude greater than what a traditional retailer observes about its wholesale suppliers.

Amazon operates one of the largest private-label programs among digital platforms, marketing products under dozens of proprietary brands—including Amazon Basics, Amazon Essentials, Solimo, Presto!, Goodthreads, Lark & Ro, and Stone & Beam—spanning consumer electronics accessories, office supplies, apparel, home furnishings, pet supplies, and grocery. Table 1 summarizes entry patterns across Amazon's product taxonomy.

***Table 1:*** Private-label entry patterns across Amazon product categories

| Group | N Categories | Avg Sale Price | New Products/Month | 3M New Products |
|---|---|---|---|---|
| Categories w PL | 1054 | 137.73 | 977.21 | 2931.42 |
| Categories w/o PL | 4074 | 162.30 | 54.41 | 161.41 |
| **Total** | **5128** | **157.25** | **244.08** | **730.75** |

*Notes:* "Categories w PL" are categories containing at least one Amazon private-label product. "Categories w/o PL" are categories with no private-label presence. U.S. marketplace. $N = 5{,}128$ total product categories.

Figure 1 plots the cumulative number of private-label and third-party seller products over time. The platform's product catalog grew rapidly after 2015, exceeding 300 million products by the early 2020s. Amazon's private-label portfolio, while orders of magnitude smaller, expanded steadily in parallel—Amazon Basics, the oldest brand, entered in 2006 and the portfolio has surpassed 40,000 cumulative products. Third-party entry has grown at a substantially faster rate, reflecting the open nature of the marketplace.

***Figure 1:*** Cumulative number of private-label and third-party seller products over time

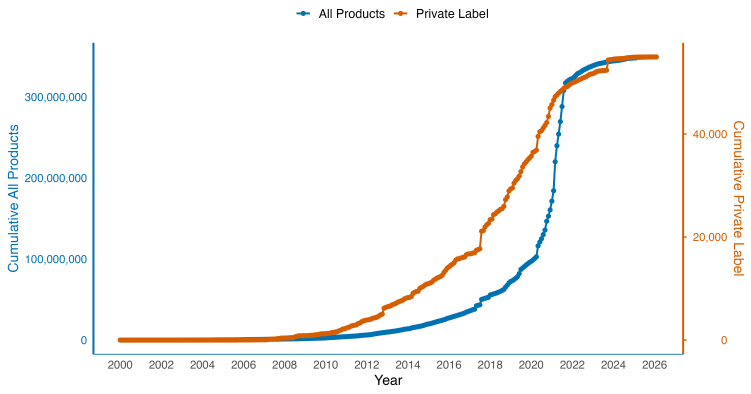
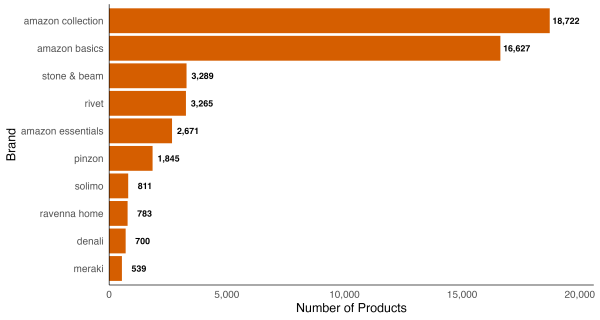



Figure 2a displays the top 10 private-label brands by product count. Amazon Basics and Amazon Collection together account for the majority of private-label products, with over 16,000 and 18,000 products respectively. The remaining brands are considerably smaller, with Stone & Beam, Rivet, and Amazon Essentials each contributing between 2,600 and 3,300 products. Figure 2b shows how the top six brands distribute their products across categories. Amazon Collection concentrates in jewelry and accessories,

Amazon Basics focuses on bedding and home goods, and Stone & Beam specializes in furniture. This category specialization suggests that Amazon tailors each brand identity to a specific market segment rather than deploying a single umbrella brand across all categories.

***Figure 2:*** Amazon private-label brand composition

***(a)*** Top 10 brands by product count




***(b)*** Category distribution for top six brands




Private labels appear in 1,054 out of 5,128 product categories. Within those categories, the average product price is $137.73, compared to $162.30 in categories without private labels. Categories with private-label presence also exhibit substantially higher entry rates: 977 new products per month versus 54 in non-private-label categories. The platform concentrates its private-label offerings in categories characterized by high search volume, large numbers of existing sellers, and relatively undifferentiated products.

The U.S. House Committee on the Judiciary's 2020 investigation raised concerns that Amazon may use proprietary marketplace data to inform the design of its private-label offerings (U.S. House Committee on the Judiciary, 2020). Additional concerns include Amazon's ability to influence product visibility through the Buy Box algorithm and search rankings (Farronato et al., 2022). These concerns are not unique to Amazon—Apple's App Store has faced similar scrutiny for replicating popular third-party apps. The European Union's Digital Markets Act restricts "gatekeepers" from using non-publicly available business-user data to compete against those users; proposed U.S. legislation would impose similar restrictions. Our empirical analysis documents the

extent of product imitation and its consequences for consumers and third-party sellers, providing evidence that is directly relevant to these regulatory discussions.

# 3 Data and Similarity Measures

## 3.1 Data Collection and Sample

We construct a novel panel dataset by scraping product pages from Amazon's four largest marketplaces: the U.S., Germany, the U.K., and Canada. Our data collection covers 2018 to the present and encompasses over 770 million product pages: approximately 350 million from the U.S., 226 million from Canada, 105 million from Germany, and 91 million from the U.K., spanning both Amazon's private-label brands and third-party sellers across the full product taxonomy.

For each product, we observe the product title, textual description, product images, sale price, manufacturer name, star rating, the number and distribution of reviews by star level, customer questions and answers, the product's position within Amazon's hierarchical category taxonomy (Category 1, Category 2, and Category 3, from broadest to most granular), sales rank within each level, image count, Buy Box status, and fulfillment method. For 162 million products, we observe the exact introduction date; for the remainder, we approximate entry using the earlier of the first customer review date or the first scraping date.

Our scraping infrastructure visits each product page approximately once per month, generating a panel that tracks changes in prices, ratings, review accumulation, sales ranks, and product availability over time. We identify Amazon's private-label products using a comprehensive list of Amazon-owned brands; all remaining products are classified as third-party.

For the imitation analysis, we focus on 14 product categories with substantial private-label presence: pet supplies (dog beds and mats), office chairs, cat condos and

cat trees, dinnerware sets, dutch ovens, bookcases, desks, suitcases, baby blankets, bath towels, accent chairs, end tables, ottomans, and food storage containers. Computing pairwise visual similarity across the full catalog of hundreds of millions of products is computationally prohibitive, requiring us to focus on a subset of categories. We select categories where products exhibit meaningful visual heterogeneity—distinct shapes, colors, and design features that allow our similarity model to discriminate between products. Categories such as cables, adapters, or standardized electronic accessories, where nearly all products look alike regardless of manufacturer, would yield uninformative similarity scores. Within these 14 categories, our sample contains 1,673 Amazon private-label products and over 6 million third-party products.

An important note on sample scope: the imitation analysis (Section 4), consumer outcomes (Section 5), and product-level regressions (Section 6) are conducted using the U.S. marketplace only, as this is the largest market with the most comprehensive data coverage. The category-level difference-in-differences analysis (Section 7) is estimated separately for each of the four marketplaces—the U.S., Germany, the U.K., and Canada—to assess the external validity of our findings.

## 3.2 Visual Similarity Measurement

To measure visual resemblance between products, we train a Siamese neural network—a model designed for learning similarity metrics (Bromley et al., 1993; Chopra et al., 2005). The network consists of two identical weight-sharing subnetworks, each mapping a product image to a $d = 512$-dimensional embedding. Let $f_\theta(\cdot)$ denote the embedding function. The training objective minimizes:

$$L(\theta) = \sum_{x_a, x_b} \underbrace{L_b(x_a, x_b)}_{\substack{\text{Penalty if similar images} \\ \text{are far apart}}} + \sum_{x_b, x_c} \underbrace{L_c(x_b, x_c)}_{\substack{\text{Penalty if dissimilar images} \\ \text{are nearby}}} \tag{1}$$

where $L_b$ is a contrastive loss that penalizes large distances between embeddings of visually similar products, and $L_c$ penalizes small distances between dissimilar products. We construct training pairs using products sharing the same manufacturer and narrowest category (similar) and products from different categories (dissimilar).

We measure similarity between products $i$ and $j$ as:

$$\text{sim}(i,j) = \frac{f_\theta(x_i) \cdot f_\theta(x_j)}{\|f_\theta(x_i)\| \, \|f_\theta(x_j)\|} \tag{2}$$

A cosine similarity of 0.9 or above corresponds to near-identical appearance—for example, office chairs of the same design differing only in branding. We compute cosine similarity for each private-label product against all third-party products within the same category. We complement visual similarity with textual similarity based on transformer-based sentence embeddings of product titles and descriptions.

### 3.3 Defining Imitation and Comparison Groups

We classify a third-party product as "potentially imitated" if there exists a private-label product in the same category with visual cosine similarity $\geq 0.9$ that was introduced after the third-party product. Under this definition, 1,222 out of 1,673 private-label products (73%) have at least one highly similar third-party predecessor. Appendix Table A3 shows that results are robust to alternative thresholds: 84% of private-label products have a match at 0.85, and 41% at the more stringent 0.95 threshold.

We emphasize that high visual similarity is consistent with deliberate imitation but also with convergent design in categories where functional constraints limit the space of viable designs. Our analysis of imitated versus non-imitated product characteristics, and our comparison with third-party imitation patterns, helps distinguish between these interpretations.

For causal estimation, we match each imitated third-party product to non-imitated

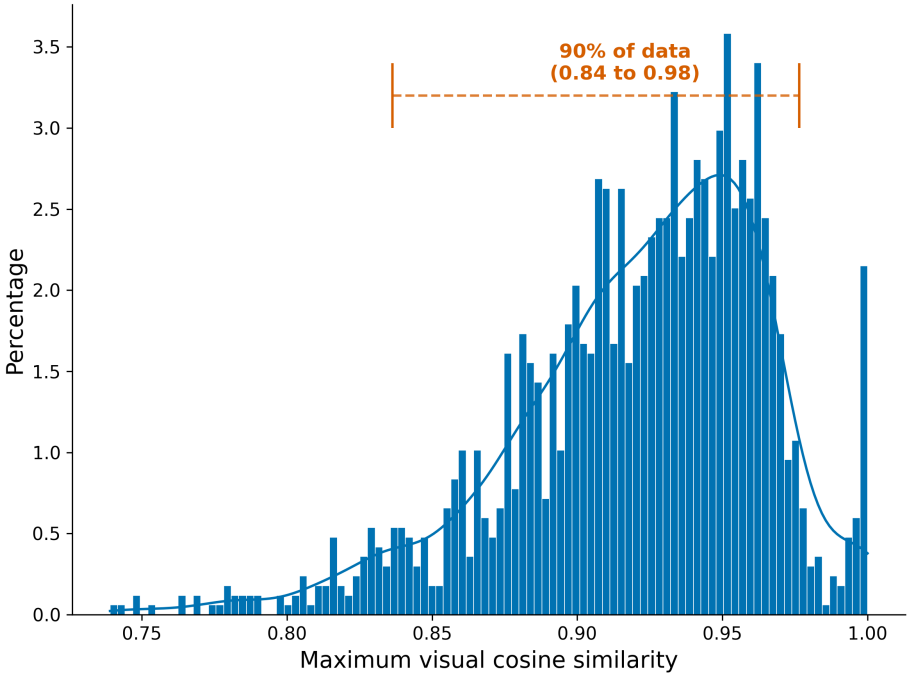
products within the same category with high similarity to the treated product ($\text{sim}(i,j) \geq 0.85$) but low similarity to any private label ($\max_k \text{sim}(j,k) < 0.7$). As a robustness check, we also construct comparison groups from non-U.S. Amazon marketplaces where private labels have not entered the relevant categories.

# 4 Evidence on Product Imitation

## 4.1 Prevalence of Visual Similarity

For each of the 1,673 private-label products in our focal categories, we compute the maximum visual cosine similarity relative to any pre-existing third-party product. Figure 3 plots this distribution.

***Figure 3:*** Distribution of maximum visual cosine similarity between each private-label product and pre-existing third-party products




Out of 1,673 private-label products, 1,222 (73%) have at least one pre-existing match with cosine similarity of 0.9 or above. The distribution is notably right-skewed, with

a pronounced mode near 0.93 and 90% of observations falling between 0.84 and 0.98. The median maximum similarity exceeds 0.92, and even the 25th percentile exceeds 0.85. Only 9% of private-label products have a maximum similarity below 0.80, indicating that virtually all private-label entries bear meaningful visual resemblance to at least one pre-existing third-party product.

At 0.9 and above, products share the same silhouette, color palette, and design elements, differing only in branding. The pervasiveness of high-similarity matches across categories—office chairs, dutch ovens, dinnerware, end tables, suitcases, dog beds—indicates a broad strategy of introducing private labels that closely resemble successful existing offerings. Figures 4a and 4b display representative product pairs; additional examples appear in Appendix B.1. In the dog beds example (Figure 4a), Amazon Basics offers a product at $37.98 that closely resembles a third-party product priced at $29.99, with a cosine similarity of 0.96. In the dutch ovens example (Figure 4b), the Amazon Basics product ($46.28) is visually near-identical to a third-party product priced at $422.99, achieving a perfect cosine similarity of 1.00—yet the private-label version is priced at roughly one-tenth of the original. These examples illustrate two distinct patterns: in some cases Amazon undercuts modestly priced competitors, while in others it offers dramatically cheaper alternatives to premium products.

## 4.2 Characteristics of Potentially Imitated Products

If private-label entry were driven purely by convergent design, imitated and non-imitated third-party products would be observationally similar along dimensions unrelated to visual appearance. Table 2 reports mean comparisons between the two groups.

Imitated products are substantially cheaper ($71.7 versus $112.1), consistent with Amazon targeting mass-market price points. They have significantly more reviews (918 versus 666), with the difference driven by four- and five-star reviews (35 and 177 more on average, respectively)—signals of commercial success that are publicly observable but

***Figure 4:*** Examples of visual similarity between Amazon private-label products and third-party predecessors

***(a)*** Dog beds (cosine similarity: 0.96)







Cosine similarity: 0.96

***(b)*** Dutch ovens (cosine similarity: 1.00)







Cosine similarity: 1.00

***Table 2:*** Characteristics of potentially imitated vs. non-imitated third-party products

| | Mean Imitated | Mean Non-Imitated | Difference | $p$-value | |
|---|---|---|---|---|---|
| Sale Price ($) | 71.7 | 112.1 | −40.4 | 0.000 | *** |
| Star Rating | 4.1 | 4.1 | −0.0 | 0.000 | *** |
| Review Count | 918.2 | 666.4 | 251.8 | 0.000 | *** |
| Image Count | 3.8 | 5.1 | −1.3 | 0.000 | *** |
| Buy Box Rank | 1.0 | 1.0 | 0.0 | 0.000 | *** |
| One-Star Reviews | 43.8 | 33.0 | 10.8 | 0.400 | |
| Two-Star Reviews | 29.5 | 20.4 | 9.0 | 0.300 | |
| Three-Star Reviews | 62.7 | 42.7 | 20.0 | 0.010 | ** |
| Four-Star Reviews | 119.8 | 84.7 | 35.1 | 0.003 | *** |
| Five-Star Reviews | 662.4 | 485.6 | 176.7 | 0.000 | *** |
| $N_{\text{imitated}} = 1{,}222$ products; | | $N_{\text{non-imitated}} = 5{,}506{,}926$ products; | | 14 categories | |

*Notes:* Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Imitated products are defined as third-party products with visual cosine similarity $\geq 0.9$ relative to at least one subsequently introduced Amazon private-label product. U.S. marketplace, 14 focal categories. The star rating difference rounds to −0.0 but is statistically significant due to the large sample size; the unrounded difference is −0.04 stars.

may also correlate with internal metrics available only to the platform. Table 3 extends this analysis to sales rank and seller characteristics. Imitated products hold significantly better (lower) sales ranks at every taxonomy level. At the broadest level (Category 1), imitated products rank approximately 187,000 positions higher; at the intermediate level (Category 2), 8,000 positions higher; and at the most granular level (Category 3), 22,000 positions higher.

***Table 3:*** Sales ranks and seller characteristics: imitated vs. non-imitated products

| | Mean Imitated | Mean Non-Imitated | Difference | $p$-value | |
|---|---|---|---|---|---|
| *Panel A: Sales Rank* | | | | | |
| Category 1 Rank | 3,039,457 | 3,225,970 | −186,513 | 0.000 | *** |
| Category 2 Rank | 7,887 | 16,005 | −8,119 | 0.000 | *** |
| Category 3 Rank | 21,183 | 43,165 | −21,982 | 0.000 | *** |
| *Panel B: Seller Characteristics* | | | | | |
| Number of Products | 5,057 | 1,531 | 3,526 | 0.000 | *** |
| Seller Tenure (days) | 3,703 | 2,968 | 735 | 0.000 | *** |
| $N_{\text{imitated}} = 1{,}222$ products; $N_{\text{non-imitated}} = 5{,}506{,}926$ products; 14 categories | | | | | |

*Notes:* Lower sales rank indicates higher sales. Seller tenure is measured as the number of days between the seller's first product listing and the date of the private-label entry. U.S. marketplace, 14 focal categories. Same sample as Table 2. *** $p < 0.01$.

Imitated products are sold by larger sellers (5,057 versus 1,531 products) who have been on the platform roughly two years longer (3,703 versus 2,968 days). Taken together, Tables 2 and 3 show that Amazon's private-label entry disproportionately targets cheaper, more popular, better-ranked products sold by larger, more experienced sellers—a systematic pattern difficult to reconcile with convergent design alone, though we cannot distinguish whether this targeting reflects proprietary data or superior market research capabilities.

## 4.3 Timing of Imitation

On average, highly similar private-label products enter 960 days (approximately 2.6 years) after the corresponding third-party product. Products sold by smaller-revenue

sellers are imitated after an average of 1,896 days (5.2 years), while products sold by larger-revenue sellers are imitated after 1,257 days (3.4 years). This gap is consistent with the platform accelerating entry for products that demonstrate commercial viability more quickly, though it could also reflect the time required for larger-revenue products to attract Amazon's attention through publicly observable signals such as reviews and best-seller rankings.

This finding has direct implications for regulatory design. Duration-based data walls may be less effective than intended if the platform's imitation lag already exceeds the proposed restriction period, particularly for smaller sellers.

## 4.4 Third-Party Imitation Patterns

Third-party sellers routinely introduce products that closely resemble existing offerings from other third-party sellers. However, the characteristics of products targeted by third-party imitators differ systematically from those targeted by Amazon.

***Table 4:*** Third-party imitation: characteristics of potentially imitated vs. non-imitated products

| | Mean Imitated | Mean Non-Imitated | Difference | $p$-value | |
|---|---|---|---|---|---|
| Sale Price ($) | 51.0 | 58.9 | −7.9 | 0.000 | *** |
| Star Rating | 4.1 | 4.1 | −0.1 | 0.000 | *** |
| Review Count | 1,650.8 | 1,320.6 | 330.3 | 0.400 | |
| Image Count | 5.4 | 5.2 | 0.2 | 0.000 | *** |
| Buy Box Rank | 1.0 | 1.0 | 0.0 | 0.000 | *** |
| One-Star Reviews | 109.9 | 115.6 | −5.7 | 0.900 | |
| Two-Star Reviews | 69.8 | 72.3 | −2.6 | 1.000 | |
| Three-Star Reviews | 140.2 | 135.6 | 4.6 | 0.900 | |
| Four-Star Reviews | 247.9 | 245.1 | 2.8 | 1.000 | |
| Five-Star Reviews | 1,431.6 | 1,407.6 | 23.9 | 0.900 | |

*Notes:* Imitated products are defined as third-party products with visual cosine similarity $\geq 0.9$ relative to at least one subsequently introduced third-party product (excluding Amazon private labels). U.S. marketplace, 14 focal categories, $N \approx 6$ million third-party products. The sample differs from Tables 2–3 because the comparison here is among third-party products only (Amazon private labels excluded from the imitator definition). Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The price gap between potentially imitated and non-imitated products is much

smaller for third-party imitation ($7.9 versus $40.4 for Amazon). The review count difference is statistically insignificant for third-party imitation, whereas it is large and highly significant for Amazon. The five-star review surplus that distinguishes Amazon's targets is absent for third-party imitators. Third-party imitators also show weaker selection on granular sales rank (see Appendix Table A4).

Two interpretations of these asymmetries merit consideration. First, Amazon's systematic selection on sales rank, review volume, and seller size could reflect access to proprietary marketplace data—such as detailed conversion rates, search analytics, and inventory turnover—unavailable to ordinary sellers. Second, the differences may reflect distinct competitive strategies rather than differential information: third-party imitators may enter at lower price points to undercut incumbents, while Amazon leverages its brand recognition and distribution infrastructure to target the most commercially established products. Our evidence cannot definitively distinguish between these channels. The systematic selection on granular sales rank—a dimension more closely tied to internal platform data than to publicly observable signals—is suggestive, but we acknowledge that a well-resourced firm could plausibly replicate this targeting through careful analysis of public information alone.

# 5 Consumer Outcomes

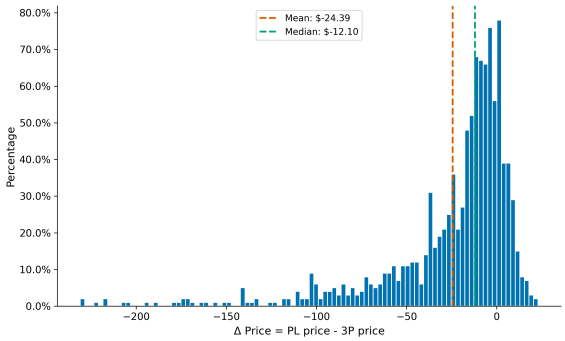
The comparisons in this section are descriptive—they characterize what consumers gain access to when Amazon introduces a private label, but do not identify a causal effect on aggregate consumer surplus. Note that the comparisons here differ from those in Section 4: the imitation analysis compared *third-party* products that were versus were not targeted by Amazon, while this section compares *private-label products directly to their closest third-party matches*.

## 5.1 Price Differences

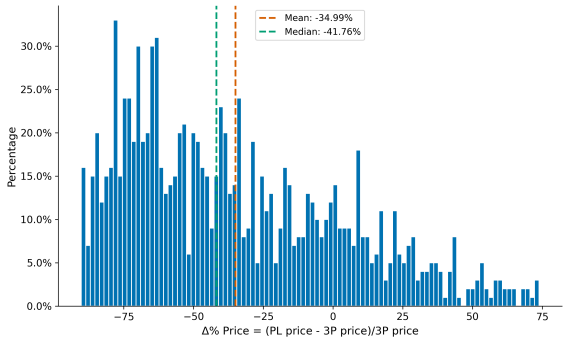
For each private-label product with a highly similar third-party predecessor (cosine similarity $\geq$ 0.9), we compute the price gap at the time of private-label introduction. Figure 5a plots the distribution of $\Delta\text{Price} = P^{PL} - P^{3P}$.

***Figure 5:*** Price differences between private-label products and their most similar third-party predecessors

***(a)*** Absolute differences ($)




***(b)*** Percentage differences




*Notes:* Sample restricted to the U.S. marketplace. Each observation is a private-label product paired with its most similar third-party predecessor (cosine similarity $\geq$ 0.9).

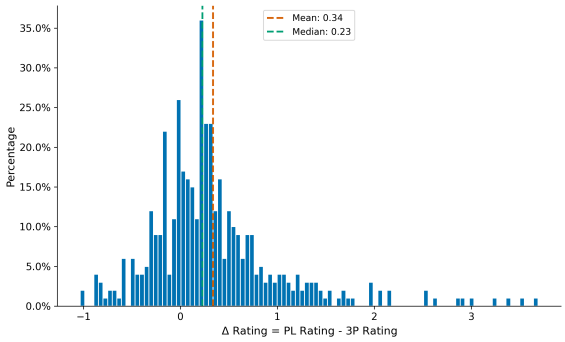
On average, Amazon's private-label products are $12 cheaper than their closest third-party matches, with the vast majority priced below their counterparts (Figure 5a). The distribution is sharply left-skewed: most price differences are concentrated between $-$\$50 and \$0, with a small number of product pairs exhibiting gaps exceeding $-$\$100, typically in higher-priced categories such as furniture and kitchen equipment. In relative terms, the median private-label product is approximately 42% cheaper (Figure 5b), and fewer than 15% of private-label products are priced above their most similar third-party predecessor. A consumer who switches to the near-identical private-label alternative saves more than two-fifths of the purchase price on average.

## 5.2 Ratings and Reviews

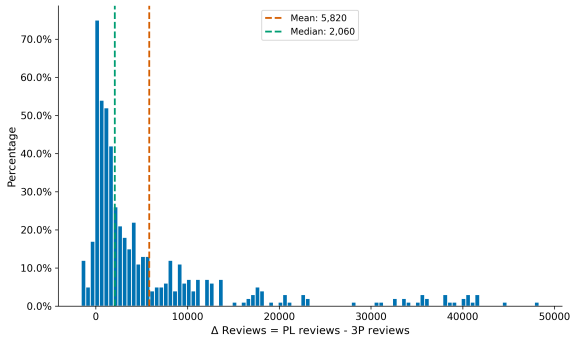
Figure 6a plots rating differences between private-label products and their most similar third-party predecessors. Private-label products carry higher average star ratings, typically by 0.1 to 0.3 stars. Even small rating differences have meaningful demand implications: Magnusson (2022) estimates that a half-star increase on Amazon is associated with a 5% increase in demand.

***Figure 6:*** Rating and review differences between private-label products and their most similar third-party predecessors

***(a)*** Star rating differences




***(b)*** Review count differences




*Notes:* Sample restricted to the U.S. marketplace. Each observation is a private-label product paired with its most similar third-party predecessor (cosine similarity $\geq 0.9$).

Figure 6a shows that the majority of private-label products carry ratings 0.1 to 0.5 stars above their third-party counterparts, with relatively few cases of lower ratings. The mean rating advantage is 0.34 stars and the median is 0.23 stars. Despite entering the market an average of 774 days after their most similar third-party predecessor (compared to 960 days across all similarity-matched pairs in Section 4.3; the shorter lag here reflects the restriction to the highest-similarity matches used in the consumer comparison), Amazon's private-label products also accumulate far more customer reviews (Figure 6b). The review count distribution is heavily right-skewed, with the mean difference exceeding 5,800 reviews and the median exceeding 2,000. Table 5 decomposes this advantage

by star rating: the surplus is overwhelmingly concentrated in five-star reviews (mean: 4,152 more; median: 1,431), while the one-star surplus is much smaller (mean: 157; median: 50). This pattern is consistent with private-label products achieving at least comparable observed customer ratings, though we note that ratings may also reflect algorithmic placement, review program participation, and selection into high-potential products rather than intrinsic quality alone.

***Table 5:*** Review count differences by star rating

| Star Rating | Mean Difference | Median Difference |
|---|---|---|
| 1-Star | 156.9 | 50.0 |
| 2-Star | 125.8 | 34.0 |
| 3-Star | 321.5 | 95.0 |
| 4-Star | 740.5 | 254.0 |
| 5-Star | 4,152.2 | 1,431.0 |

*Notes:* Differences computed as $N_s^{PL} - N_s^{3P}$ for each star level $s \in \{1, 2, 3, 4, 5\}$. Positive values indicate the private-label product has more reviews. Based on 1,222 matched PL–3P product pairs. U.S. marketplace.

The large review volume advantage implies substantially higher sales volumes, higher review propensity, or both. Contributing factors may include lower prices, prominent algorithmic placement (Farronato et al., 2022), the "Amazon's Choice" badge, consumer trust in Amazon-branded products, selection bias (Amazon may launch private labels only for products it expects to perform well), and Amazon's Vine early reviewer program. The short-run consumer implications appear favorable, but the long-run welfare consequences depend on how private-label entry affects third-party sellers' incentives—a question we address next.

# 6 Product-Level Effects

We now turn to the causal effects of private-label entry on third-party sellers. We begin with product-level estimates that isolate the effect on individual products directly

targeted by Amazon's private labels, before examining broader category-level outcomes in the next section.

For each third-party product matched to an Amazon private label (cosine similarity $\geq 0.9$), we define treatment as the month in which the similar private-label product first appears. The control group consists of third-party products in the same category that are similar to the treated product but dissimilar to any private label (see Section 3.3 for matching details). We estimate product-level difference-in-differences regressions including product fixed effects—which absorb all time-invariant product characteristics—and group-year-by-week fixed effects, which flexibly control for common shocks within cohorts of products treated at the same time. Table 6 reports the results.

***Table 6:*** Product-level effects of Amazon private-label entry on matched third-party products

| | Log(Sales Rank) | Log(Sales Rank, Granular) | Log(Price) | Discontinued |
|---|---|---|---|---|
| After Amazon Entry | 0.045*** | 0.040*** | −0.011*** | 0.003** |
| | (0.010) | (0.009) | (0.001) | (0.001) |
| Product FE | ✓ | ✓ | ✓ | ✓ |
| Group-Year × Week FE | ✓ | ✓ | ✓ | ✓ |
| # Matched 3P Products | 1,831 | 1,831 | 4,141 | 4,141 |
| Observations | 108,446 | 108,444 | 292,897 | 292,897 |

*Notes:* Each column reports the coefficient on a post-treatment indicator from a product-level difference-in-differences regression. The sample consists of third-party products in the U.S. marketplace (14 focal categories) with visual cosine similarity $\geq 0.9$ to a subsequently introduced Amazon private-label product. The number of matched products varies across specifications because sales rank data is available for fewer products than price and discontinuation data. Standard errors clustered at the product level in parentheses. ** $p < 0.05$, *** $p < 0.01$.

After Amazon introduces a visually similar private-label product, matched third-party products experience a statistically significant worsening of sales rank: 4.5% at the broad taxonomy level and 4.0% at the granular level. Because sales rank is a monotonic function of recent sales velocity, the rank deterioration implies a decline in relative sales performance. However, the mapping from rank to sales is nonlinear, so these coefficients should be interpreted as directional evidence of competitive harm

rather than as precise estimates of sales losses. Matched products also reduce their prices by 1.1%, consistent with competitive pressure from the lower-priced private-label entrant. This price decline is modest relative to the 41% average price gap documented in Section 5, suggesting that incumbents adjust prices only partially rather than fully matching the private-label price.

The probability of product discontinuation increases by 0.3 percentage points. While the magnitude appears small in absolute terms, the baseline discontinuation rate in our sample is low, making this a proportionally meaningful increase. The discontinuation result also provides direct evidence that private-label entry contributes to product exit, complementing the aggregate seller decline we document at the category level.

These product-level results demonstrate that the competitive effects of private-label entry extend directly to the performance and survival of specific targeted incumbents. In the next section, we complement these estimates with category-level analysis to assess the broader competitive restructuring induced by private-label entry.

# 7 Category-Level Effects

## 7.1 Empirical Strategy

We estimate the effect of Amazon's private-label entry on category-level outcomes using a staggered difference-in-differences design. The unit of observation is a product category–month. Treatment is defined as the first appearance of an Amazon private-label product in a category. Categories that never receive a private label serve as the comparison group. Because treatment is staggered—categories are first treated in different months—we implement the Callaway and Sant'Anna (2021) estimator, which avoids the well-known biases of two-way fixed effects under heterogeneous treatment

effects:

$$ATT(g,t) = \mathbb{E}\big[Y_t(g) - Y_t(0) \mid G = g\big] - \mathbb{E}\big[Y_t(g) - Y_t(0) \mid G = \infty\big] \tag{3}$$

where $g$ indexes the cohort (month of first treatment), $G = \infty$ denotes the never-treated group, and the overall $ATT$ aggregates across cohorts and post-treatment periods. Standard errors are clustered at the category level. We estimate the model separately for each of four Amazon marketplaces: the U.S., Germany, the U.K., and Canada.

We report only outcomes for which the event study exhibits clean pre-treatment dynamics—coefficients centered on zero with no visible trend. Outcomes that fail this criterion (minimum price, out-of-stock rates, sales ranks) are reported in the appendix.

## 7.2 Seller Exit

Private-label entry is associated with a large, sustained relative decline in the number of products operating in a category. Figure 7 plots the event study coefficients for log total sellers in each market. In all four panels, pre-treatment estimates are tightly centered on zero, and a monotonic decline begins precisely at the treatment date. The U.S. aggregate ATT is $-0.90$ (SE $= 0.075$). The decline is even steeper in Germany ($-1.65$), the U.K. ($-0.99$), and Canada ($-0.97$), all significant at the 0.1% level (Table 7).

Interpreting the magnitude requires care, for two reasons. First, the baseline sample includes many small and near-empty categories. When we restrict to categories with at least 50 products at baseline, the U.S. ATT attenuates to $-0.47$ (SE $= 0.062$), implying that treated categories have 37% fewer products than they would have absent Amazon's entry—still large and highly significant, but substantially smaller than the full-sample estimate. The estimates for Germany ($-1.32$), the U.K. ($-1.32$), and Canada ($-1.08$) are robust to this restriction (Appendix E). We report the full-sample estimates in Table 7 for comparability across outcomes, but the attenuated U.S. estimate likely provides a more

reliable benchmark.

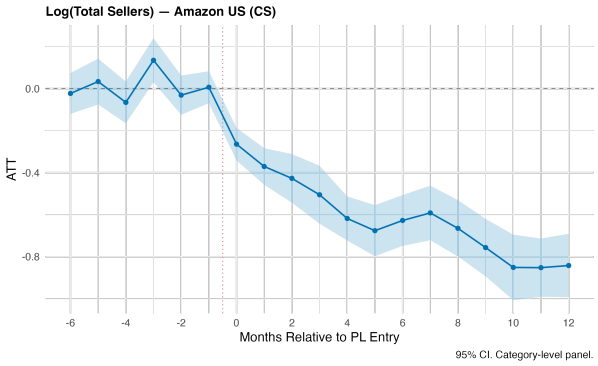
Second, the ATT on the log scale captures how much lower the product count in treated categories is compared to what it would have been absent Amazon's entry, which on a growing marketplace combines two margins: increased product exit and suppressed new entry. In Appendix E, we decompose these margins using both raw counts and entry/exit rates normalized by baseline category size. Estimating the CS model on entry and exit rates yields a striking finding: both the entry rate and the exit rate *increase* significantly after private-label entry (U.S. entry rate: $+0.24$; exit rate: $+0.34$; both significant at the 1% level). The exit rate rises more than the entry rate, producing a significantly negative net change rate ($-0.10$). This pattern holds across all four markets. Private-label entry thus increases competitive churn—more products entering and more products exiting—with exits outpacing entries. The net decline is not driven by entry suppression alone but rather by an acceleration of competitive turnover that disproportionately favors exit.

## 7.3 Price Compression

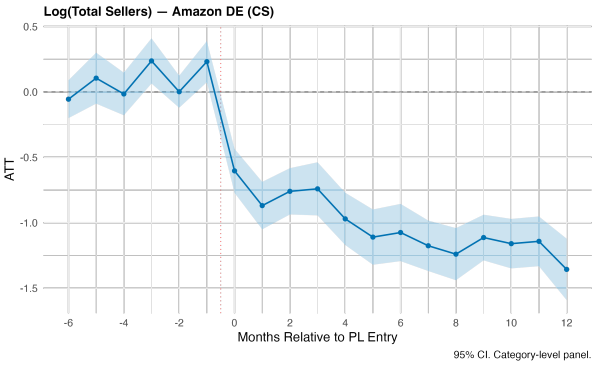
Private-label entry compresses the upper tail of the price distribution within a category—that is, the highest prices observed in a category decline after Amazon enters, either because high-priced sellers reduce their prices in response to competition or because they exit the category entirely. We measure this by tracking the log of the maximum price in each category over time. Figure 8 plots the event study for this outcome. In the U.S., pre-treatment coefficients hover around zero and then drop sharply at treatment, stabilizing at an approximately 26–33% decline. The aggregate ATT is $-0.49$ (SE $= 0.064$), corresponding to a 39% decline in the maximum category price. The pattern is consistent across markets: the U.K. ATT is $-0.56$ (43% decline), with larger magnitudes in Germany ($-0.95$, 61%) and Canada ($-0.65$, 48%). The highest-priced products in a category either reduce their prices or exit after Amazon enters.

***Figure 7:*** Effect of private-label entry on log total sellers

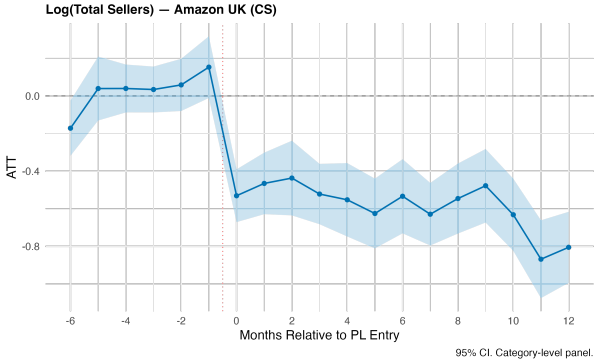
***(a)*** United States



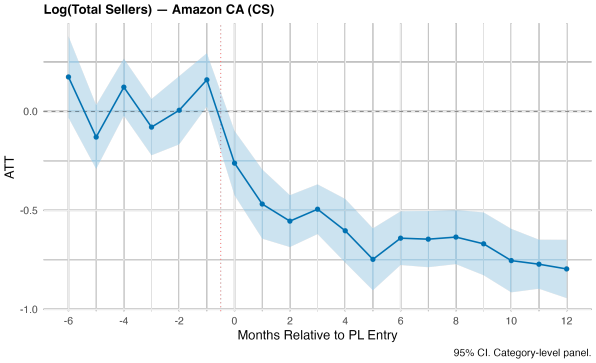

***(b)*** Germany




***(c)*** United Kingdom



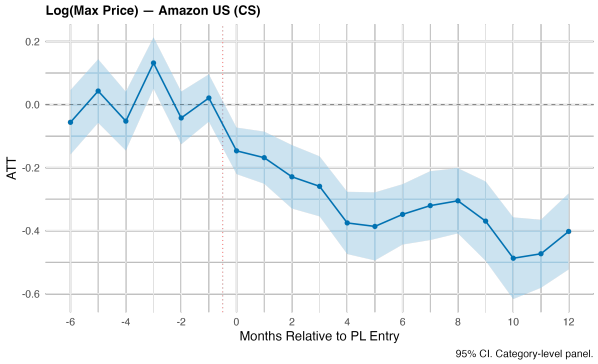

***(d)*** Canada



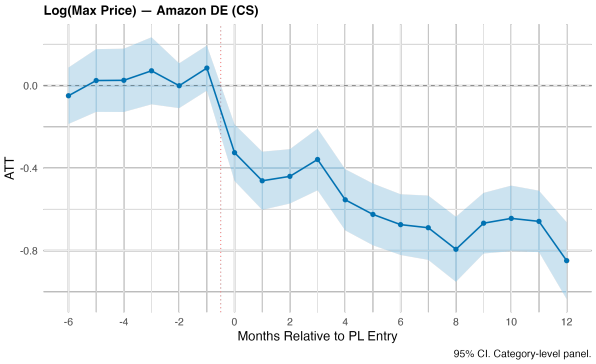

*Notes:* Each panel plots the Callaway and Sant'Anna (2021) group-time $ATT(g,t)$ estimates with 95% pointwise confidence intervals. The $x$-axis measures months relative to first private-label entry; the dashed vertical line marks treatment. Pre-treatment coefficients near zero support the parallel trends assumption.

***Figure 8:*** Effect of private-label entry on log maximum price

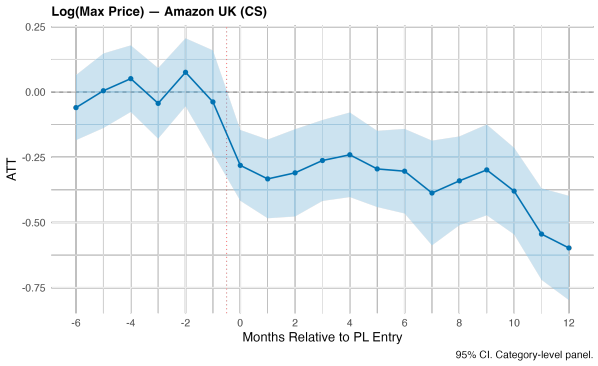
***(a)*** United States



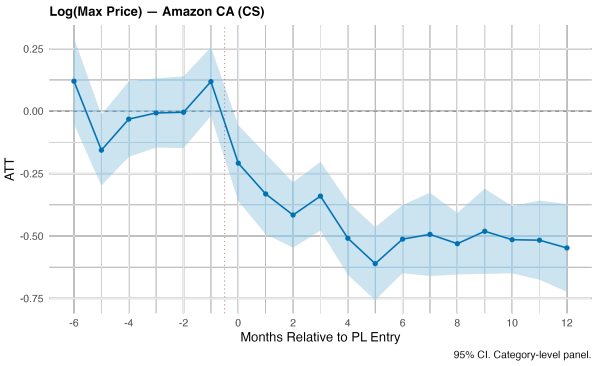

***(b)*** Germany




***(c)*** United Kingdom



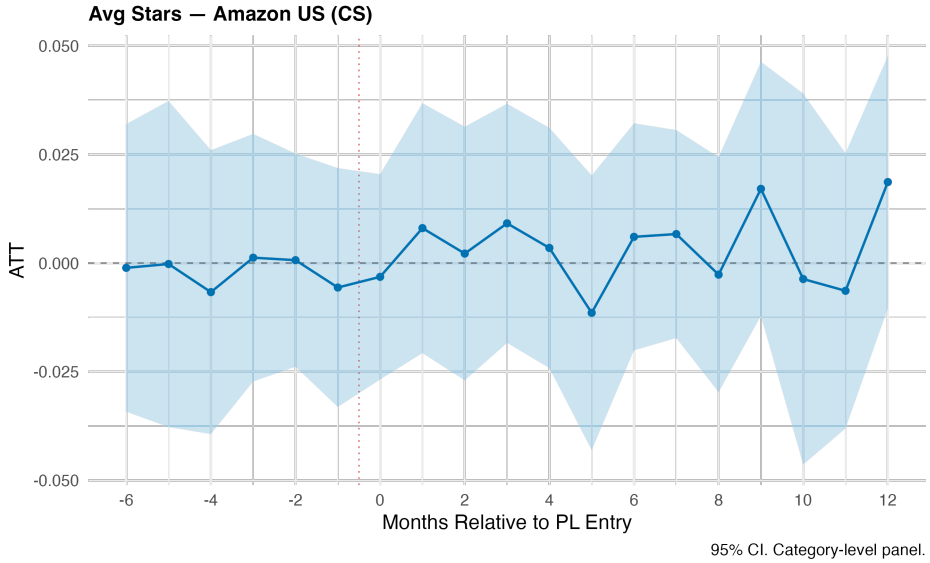

***(d)*** Canada




*Notes:* See notes to Figure 7. The maximum price in the category declines at treatment in all four markets.

An important concern is whether the price decline reflects a shift in the composition of products toward lower quality. Figure 9 shows that average star ratings are essentially unaffected by private-label entry (U.S. ATT $= -0.006$, statistically insignificant), with no pre-trend and no post-treatment movement. The price compression therefore reflects changes in pricing or seller composition, not quality degradation.

***Figure 9:*** Effect of private-label entry on average star rating (U.S.)




*Notes:* See notes to Figure 7. The null result is consistent across all four markets (Table 7).

## 7.4 Listing Quality and Reviews

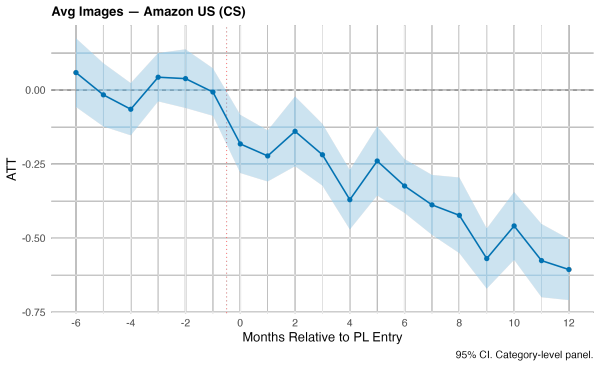
Private-label entry is followed by a decline in listing investment and a loss of accumulated reviews. Figure 10 presents the U.S. event studies for the two main content outcomes. Panel (a) shows that the average number of images per listing falls by 0.67 (SE $= 0.040$), a 19% decline relative to the baseline mean of 3.5. The event study is clean: pre-treatment coefficients are flat near zero, and the decline begins at treatment. Panel (b) shows that log total reviews falls by 0.68 (SE $= 0.093$), with pre-treatment coefficients oscillating near zero and a sustained downward shift post-treatment. Because reviews accumulate on

products over their lifetime, this decline is consistent with the exit of established products from treated categories.

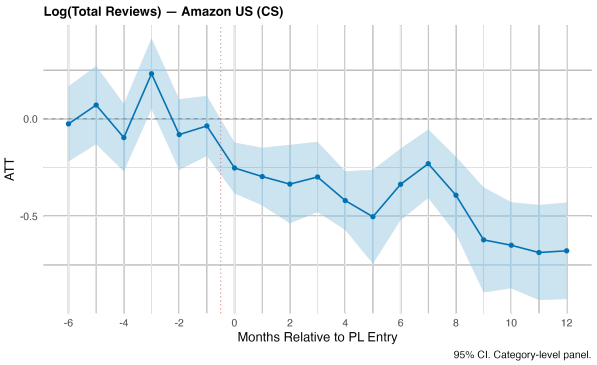
Both effects replicate across markets. Image counts fall by 0.97 in Germany, 0.85 in the U.K., and 0.69 in Canada. Log total reviews falls by 1.45 in Germany, 1.15 in the U.K., and 0.53 in Canada (all significant at the 1% level; Table 7).

***Figure 10:*** Effect of private-label entry on listing content and reviews (U.S.)

***(a)*** Average images per listing




***(b)*** Log total reviews




*Notes:* See notes to Figure 7. Both outcomes exhibit flat pre-trends and sustained post-treatment declines.

## 7.5 Summary

Table 7 collects the aggregate ATT estimates for all outcomes with clean parallel trends. Three findings stand out. First, the dominant pattern is competitive contraction—fewer sellers, declining maximum prices, and a smaller review stock. The one exception is average video count, which increases modestly in some markets, possibly reflecting a concurrent platform-wide shift toward video content rather than a treatment effect. Second, the magnitudes are economically large: treated categories having 37–57% fewer products than they would have absent Amazon's entry (in the conservative specification restricting to categories with at least 50 baseline products) and a 19–28% decline in listing images represent substantial restructuring of treated categories. Third, product quality as measured by star ratings is largely unaffected in the U.S., Germany, and the U.K.,

though Canada exhibits a small but statistically significant decline ($-0.053$), suggesting that quality effects may emerge in smaller or more concentrated markets.

***Table 7:*** Category-level effects of private-label entry: Callaway–Sant'Anna ATT

| | Amazon Marketplace | | | |
|---|---|---|---|---|
| Outcome | U.S. | DE | U.K. | CA |
| *Panel A: Competition and Prices* | | | | |
| Log(Total Sellers) | $-0.901^{***}$ | $-1.650^{***}$ | $-0.990^{***}$ | $-0.974^{***}$ |
| | (0.075) | (0.084) | (0.100) | (0.084) |
| Log(Max Price) | $-0.493^{***}$ | $-0.948^{***}$ | $-0.557^{***}$ | $-0.650^{***}$ |
| | (0.064) | (0.060) | (0.080) | (0.078) |
| *Panel B: Listing Quality* | | | | |
| Avg Images | $-0.672^{***}$ | $-0.968^{***}$ | $-0.852^{***}$ | $-0.694^{***}$ |
| | (0.040) | (0.054) | (0.066) | (0.052) |
| Log(Total Reviews) | $-0.684^{***}$ | $-1.451^{***}$ | $-1.153^{***}$ | $-0.531^{***}$ |
| | (0.093) | (0.095) | (0.128) | (0.135) |
| Avg Videos | $0.058^{***}$ | $0.009^{*}$ | $-0.005$ | $0.312^{***}$ |
| | (0.013) | (0.005) | (0.006) | (0.046) |
| *Panel C: Product Quality* | | | | |
| Avg Star Rating | $-0.006$ | $-0.016$ | $-0.014$ | $-0.053^{***}$ |
| | (0.016) | (0.028) | (0.021) | (0.017) |
| Categories (treated + control) | 8,351 | 5,561 | 5,874 | 6,382 |
| Observations | 317,281 | 208,470 | 222,513 | 242,238 |

*Notes:* Each cell reports the aggregate *ATT* from the Callaway and Sant'Anna (2021) estimator. The unit of observation is a product category–month. Treated categories received their first Amazon private-label product during 2019–2025; never-treated categories serve as the comparison group. Observation counts correspond to the competition outcomes (log total sellers); counts vary slightly across outcome groups because not all categories have data for every outcome in every month (e.g., some categories lack review or image data in certain periods). The number of categories also differs across markets because each marketplace has a distinct product taxonomy. Only outcomes satisfying a visual pre-trends check are reported; see the appendix for outcomes with problematic pre-trends. Standard errors clustered at the category level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The effects are largest in Germany, where seller exit and review losses are roughly twice the U.S. magnitudes, and smallest—but still highly significant—in Canada. The consistency of signs across four independent markets strengthens the causal interpretation and reduces concerns about market-specific confounds. Taken together,

these results indicate that private-label entry triggers a broad competitive restructuring that extends across multiple outcome dimensions and marketplace geographies.

# 8 Conclusion

This paper provides the first large-scale empirical analysis of product imitation by a dominant platform and its consequences for consumers and third-party sellers. Applying image recognition and textual similarity methods to over 770 million Amazon product pages across four marketplaces, we show that 73% of Amazon's private-label products in our focal categories are visually near-identical to pre-existing third-party products, and that the platform disproportionately targets cheaper, more popular products sold by larger, more established sellers. Amazon's imitation patterns differ qualitatively from imitation among third-party sellers, a difference that may reflect access to proprietary marketplace data or distinct competitive strategies.

The welfare consequences of private-label entry involve a clear tradeoff. Consumers gain access to products that are 41% cheaper with comparable or higher ratings. At the category level, however, private-label entry causes substantial competitive restructuring: treated categories have significantly fewer products than they would have absent Amazon's entry—driven by increased competitive churn in which both entry and exit rates rise, with exits outpacing entries—the highest prices in treated categories fall sharply (maximum price declines by approximately 39–61%), and the cumulative review stock erodes. At the product level, individual targeted third-party products experience worsening sales ranks, price declines, and higher discontinuation rates. Product quality, as measured by star ratings, is largely unaffected. These effects are remarkably consistent across all four Amazon marketplaces we study—the U.S., Germany, the U.K., and Canada—with Germany exhibiting the largest magnitudes.

These findings inform the design of platform regulation. The evidence of systematic

product imitation supports concerns raised by the U.S. House Judiciary Committee and motivates restrictions in the EU's Digital Markets Act and proposed U.S. legislation. At the same time, the substantial consumer benefits caution against eliminating private-label entry altogether. The category-level results suggest that the primary competitive harm operates through seller displacement rather than quality degradation, pointing toward regulatory frameworks that protect competitive entry incentives while preserving the consumer price benefits of private labels.

Several limitations suggest directions for future work. Our similarity-based measure cannot definitively distinguish deliberate imitation from convergent design; structural models of the platform's entry decision could provide sharper identification. While we decompose the product count decline into entry and exit rate changes (Appendix E), we cannot characterize *which* products exit—whether the displaced products are low-quality fringe offerings (suggesting efficiency-enhancing creative destruction) or established high-quality sellers (suggesting harmful displacement). Distinguishing these scenarios is critical for welfare assessment and represents an important direction for future research. Extending this framework to other dual-mode platforms—Apple's App Store, Walmart's marketplace—would establish the generality of these patterns and inform platform-neutral regulatory principles.

# References


**Anderson, Simon P. and Özlem Bedre-Defolie**, "Platform Self-Preferencing and Private Label Entry," *Working Paper*, 2025.

**Bromley, Jane, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah**, "Signature Verification Using a "Siamese" Time Delay Neural Network," in "Advances in Neural Information Processing Systems," Vol. 6 1993.

**Callaway, Brantly and Pedro H. C. Sant'Anna**, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

**Chen, Liang and Yiwei Tsai**, "Amazon's Private Label Investments," *Working Paper*, 2022.

**Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim**, "Platform Design and Private Labels," *Working Paper*, 2025.

**Chopra, Sumit, Raia Hadsell, and Yann LeCun**, "Learning a Similarity Metric Discriminatively, with Application to Face Verification," in "Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition" 2005, pp. 539–546.

**Crawford, Gregory S., Pascal Courty, Chiara Bentivogli et al.**, "Platform Self-Preferencing: Evidence from Digital Markets," *Working Paper*, 2022.

**Dubé, Jean-Pierre**, "Private Labels and Retail Competition," *Working Paper*, 2024.

**Farrell, Joseph and Michael L. Katz**, "Innovation, Rent Extraction, and Integration in Systems Markets," *Journal of Industrial Economics*, 2000, *48* (4), 413–432.

**Farronato, Chiara, Andrey Fradkin, and Alexander MacKay**, "Self-Preferencing at Amazon: Evidence from Search Rankings," *AEA Papers and Proceedings*, 2022, *112*, 239–243.

**_ , _ , and _** , "The Quality of Amazon Private Labels," *Working Paper*, 2025.

**Gutierrez, Germán**, "The Welfare Consequences of Regulating Amazon," *Working Paper*, 2021.

**Hagiu, Andrei and Daniel Spulber**, "First-Party Content and Coordination in Two-Sided Markets," *Management Science*, 2013, *59* (4), 933–949.

**_ , Bruno Jullien, and Julian Wright**, "Creating Platforms by Hosting Rivals," *Management Science*, 2015.

**_ , Tat-How Teh, and Julian Wright**, "Should Platforms Be Allowed to Sell on Their Own Marketplaces?," *RAND Journal of Economics*, 2022, *53* (2), 297–327.

**Hervas-Drane, Andres and Sandro Shelegia**, "Marketplace or Reseller?," *Management Science*, 2022, *68* (6), 4092–4107.

**Jeffries, Daniel and Pai-Ling Yin**, "Platform Entry and Innovation: Evidence from Amazon Private Labels," *Working Paper*, 2021.

**Jiang, Baojun, Kinshuk Jerath, and Kannan Srinivasan**, "Firm Strategies in the "Mid Tail" of Platform-Based Retailing," *Marketing Science*, 2011, *30* (5), 757–775.

**Lam, H. Tai**, "Private Labels on Platforms," *Working Paper*, 2024.

**Lee, Kwok Hao and Leon Musolff**, "Entry into Two-Sided Markets Shaped by Platform-Guided Search," *Working Paper*, 2021.

**Madsen, Erik and Nikhil Vellodi**, "Insider Imitation," *Working Paper*, 2022.

**Magnusson, Niklas**, "Star Ratings and Demand for Products on Amazon," *Working Paper*, 2022.

**Raval, Devesh**, "Platform Private Labels and Search Bias," *Working Paper*, 2022.

**U.S. House Committee on the Judiciary**, "Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations," Staff Report, Subcommittee on Antitrust, Commercial and Administrative Law, Washington, DC 2020.

**Zou, Tianshu and Bo Zhou**, "Platform Private Labels and Third-Party Seller Innovation," *Working Paper*, 2023.

# Appendix

# A Additional Details on Data Construction

## A.1 Data Completeness and Scraping Frequency

Our scraping infrastructure collects product pages from Amazon on a rolling basis. Due to the enormous scale of the catalog—over 770 million products across four Amazon marketplaces—we cannot visit each product page daily. On average, each product is scraped approximately once per month. The scraping frequency is higher for products in categories with active private-label presence, as these categories are prioritized in our collection schedule.

For the 162 million products with exact introduction dates recorded in Amazon's metadata, we use those dates directly. For the remaining products, we approximate the introduction date using the earlier of (i) the date of the product's first customer review and (ii) the first date on which the product appears in our scraping sample. The first-review approximation is likely accurate to within a few weeks for popular products, which tend to receive reviews shortly after launch. For less popular products, the approximation may introduce measurement error of one to several months. We assess the sensitivity of our results to this approximation by restricting the sample to products with exact introduction dates and find that our main estimates are qualitatively unchanged.

Table A1 summarizes the distribution of scraping frequency across products in our sample.

***Table A1:*** Distribution of scraping frequency

| | Mean | Median | 25th Pctl. | 75th Pctl. |
|---|---|---|---|---|
| Days between scrapes | 31.2 | 28 | 21 | 38 |
| Number of observations per product | 14.8 | 12 | 7 | 20 |

*Notes:* Statistics computed over all products in the 14 focal categories with at least two observations. U.S. marketplace. $N \approx 6$ million products.

## A.2 Private-Label Brand Identification

We identify Amazon's private-label products using a comprehensive list of Amazon-owned brands compiled from public reporting, Amazon's own disclosures, and manual verification. Table A2 lists the private-label brands in our data, ordered by the number of products in our sample.

***Table A2:*** Amazon private-label brands in our sample

| Brand Name | # Products | Primary Categories |
|---|---|---|
| Amazon Collection | 18,722 | Jewelry, accessories |
| Amazon Basics | 16,627 | Electronics accessories, office, home |
| Stone & Beam | 3,289 | Furniture, home décor |
| Rivet | 3,265 | Furniture, lighting |
| Amazon Essentials | 2,671 | Apparel, basics |
| Pinzon | 1,845 | Bedding, bath |
| Solimo | 811 | Grocery, personal care, household |
| Ravenna Home | 783 | Furniture |
| Denali | 700 | Outdoor, home |
| Meraki | 539 | Personal care |

*Notes:* Product counts reflect the total number of unique products observed across all categories in the full scraping sample, not limited to the 14 focal categories used in the imitation analysis.

## A.3 Product Taxonomy

Amazon organizes products into a hierarchical taxonomy with multiple levels of granularity. We refer to these levels as Category 1 (broadest), Category 2, and Category 3 (most granular). For example, a dog bed might be classified as: Category 1 = "Pet Supplies," Category 2 = "Dogs," Category 3 = "Beds & Furniture." Sales ranks are computed separately within each taxonomy level, so a product's Category 3 rank reflects its performance relative to other products in the same narrow segment.

We use the Category 3 classification for within-category matching and similarity computation, as it provides the most precise competitive segmentation. Category 1 and Category 2 ranks are used as additional outcome variables in our analysis to assess

whether the effects of private-label entry propagate across taxonomy levels.

# B Additional Results on Product Similarity

## B.1 Visual Similarity Examples

Figures A1–A4 display additional side-by-side comparisons of Amazon private-label products and their most similar third-party predecessors across our focal categories (see also Figures 4a and 4b in the main text). Each figure shows the private-label product, the matched third-party product, and the cosine similarity score between them. These examples illustrate the degree of visual resemblance captured by our Siamese neural network embedding model.

## B.2 Distribution of Similarity Scores over Time

An important question is whether the degree of similarity between private-label products and their third-party predecessors has changed over time. Figure A5 plots the average maximum cosine similarity for private-label products introduced in each year of our sample. The figure reveals a declining trend: more recently introduced private-label products tend to have lower maximum similarity scores relative to pre-existing third-party products. This pattern could reflect several forces, including increased regulatory scrutiny, changes in Amazon's private-label strategy, or the exhaustion of the most easily imitated product niches.

## B.3 Robustness to Alternative Similarity Thresholds

Our main analysis defines imitation using a cosine similarity threshold of 0.9. Table A3 reports the number of private-label products classified as having a highly similar third-party predecessor under alternative thresholds ranging from 0.80 to 0.95. The

***Figure A1:*** Office chairs: Amazon private-label product and most similar third-party predecessor



***Figure A2:*** Dinnerware sets: Amazon private-label product and most similar third-party predecessor

| $57.75<br>amazon basics | $138.97<br>yangyang |
|---|---|

Cosine similarity: 1.00





***Figure A3:*** End tables: Amazon private-label product and most similar third-party predecessor

| $39.74 | $204.57 |
| --- | --- |
| amazon basics | jymbk |

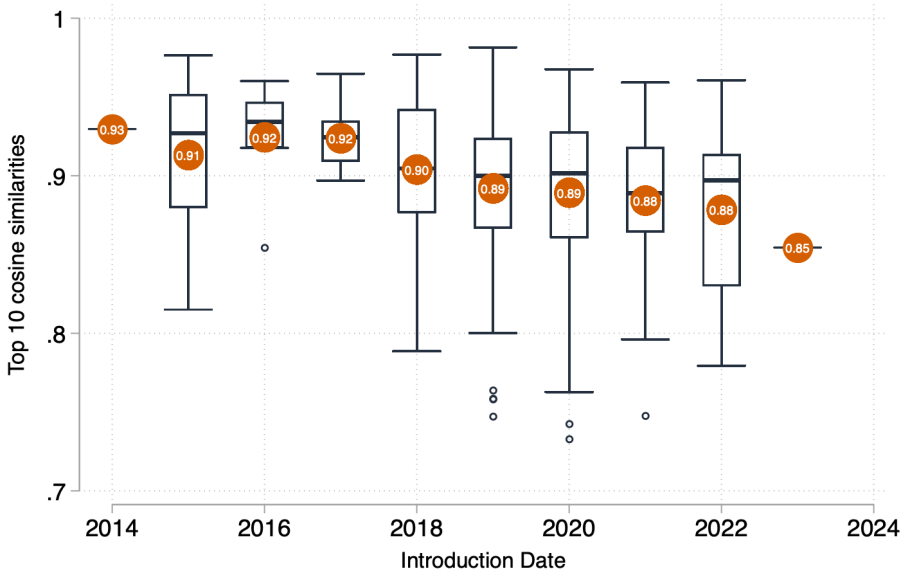
Cosine similarity: 1.00



***Figure A4:*** Suitcases: Amazon private-label product and most similar third-party predecessor



***Figure A5:*** Average maximum visual cosine similarity between Amazon private-label products and pre-existing third-party products, by year of private-label introduction




results confirm that the prevalence of high-similarity matches is not an artifact of our particular threshold choice. Even at the more stringent threshold of 0.95, over 40% of private-label products have a near-identical third-party predecessor.

***Table A3:*** Prevalence of similarity under alternative thresholds

| Similarity Threshold | # PL Products with Match | Percent of Total |
|---|---|---|
| $\geq 0.80$ | 1,521 | 90.9% |
| $\geq 0.85$ | 1,408 | 84.2% |
| $\geq 0.90$ | 1,222 | 73.1% |
| $\geq 0.95$ | 687 | 41.1% |

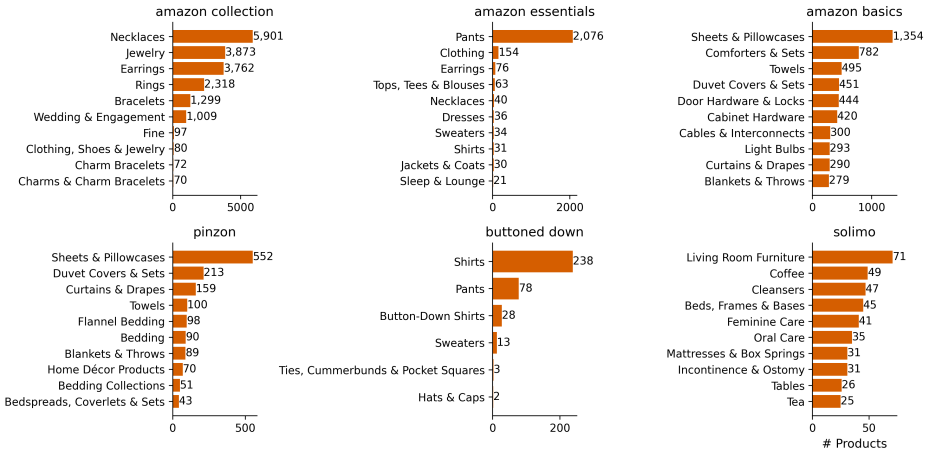
*Notes:* Total number of private-label products in the 14 focal categories is 1,673 ($N_{3P} \approx 6$ million). U.S. marketplace. A private-label product is classified as having a match if at least one pre-existing third-party product in the same category has a visual cosine similarity at or above the threshold.

# C Category-Level Descriptive Statistics

## C.1 Private-Label Entry across Categories

Figure A6 shows the distribution of private-label products across product categories. Private-label entry is concentrated in a relatively small number of categories: the top 20 categories account for over 60% of all private-label products. The most heavily targeted categories tend to be those with high search volume, large numbers of existing third-party sellers, and products that are relatively undifferentiated.

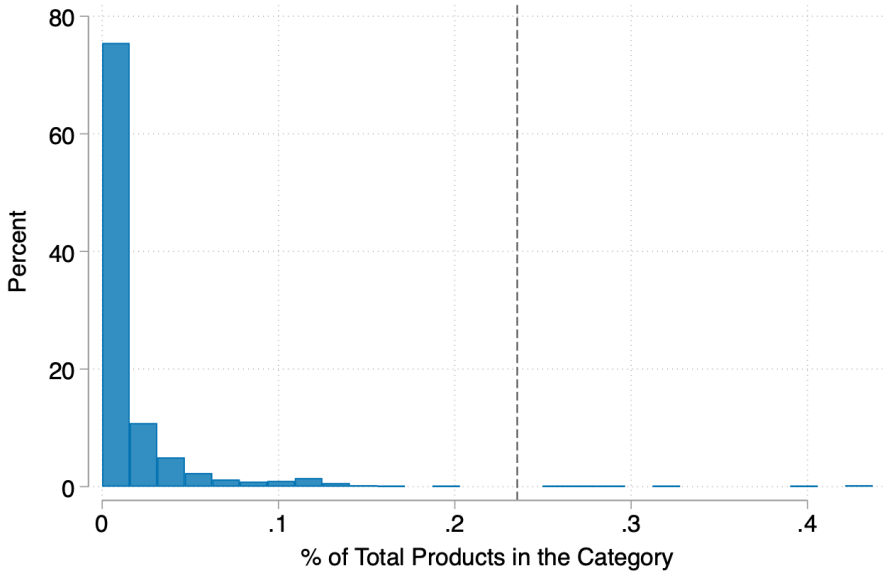
***Figure A6:*** Number of Amazon private-label products per product category




## C.2 Share of Private Labels by Category

While private labels account for only 0.01% of the overall Amazon catalog, their share varies substantially across categories. Figure A7 plots the distribution of private-label market share (measured as the fraction of products in each category that are private labels) across all categories in which at least one private-label product is present. The distribution is heavily right-skewed: in most affected categories, private labels account for less than 1% of products, but in a small number of categories the share exceeds 5%.

*Figure A7:* Distribution of private-label product share across categories




# D Additional Results on Seller Outcomes

## D.1 Third-Party-to-Third-Party Imitation: Sales Rank Comparison

Table A4 reports sales rank comparisons for third-party-to-third-party imitation, the analogue of the Amazon imitation analysis in Table 3 of the main text. While products imitated by other third-party sellers do tend to have better Category 1 and Category 2 ranks, the differences are smaller than those observed for Amazon's imitation, and the Category 3 rank difference is statistically insignificant.

*Table A4:* Third-party imitation: sales rank comparisons

| Taxonomy Level | Mean Imitated | Mean Non-Imitated | Difference | $p$-value | |
|---|---|---|---|---|---|
| Category 1 | 496,649 | 995,456 | −498,807 | 0.000 | *** |
| Category 2 | 7,342 | 14,134 | −6,792 | 0.003 | *** |
| Category 3 | 23,428 | 32,910 | −9,482 | 0.600 | |

*Notes:* Imitated products are defined as third-party products with visual cosine similarity $\geq 0.9$ relative to at least one subsequently introduced third-party product (excluding Amazon private labels). Lower sales rank indicates higher sales. U.S. marketplace, 14 focal categories. $N \approx 6$ million third-party products. Stars indicate statistical significance: ** $p < 0.05$, *** $p < 0.01$.

The insignificance of the Category 3 rank difference is particularly informative. Category 3 represents the most granular level of Amazon's taxonomy and therefore the most precise measure of a product's competitive position within its niche. The fact that third-party imitators do not systematically target products with superior granular rankings—while Amazon does—is suggestive, though as discussed in the main text, we cannot definitively distinguish whether this reflects proprietary data access or Amazon's superior market research capabilities.

## D.2 Imitation Timing: Heterogeneity by Seller Size

Figure A8 disaggregates the distribution of imitation lags by seller revenue quintile. Smaller sellers (lower revenue quintiles) experience substantially longer lags between their product's introduction and the entry of a similar Amazon private label. This heterogeneity has direct implications for the design of duration-based data-use restrictions: a five-year data wall would bind for the average large seller (imitated after 3.4 years) but would already have expired for the average small seller (imitated after 5.2 years).

## D.3 Distribution of Imitation Lags

Figure A9 plots the full distribution of the number of days between the introduction of a third-party product and the entry of a highly similar (cosine similarity $\geq 0.9$) Amazon private-label product. The distribution is right-skewed with a long tail: the median lag is approximately 800 days, and the mean is 960 days. A non-trivial fraction of imitation events occur more than five years after the original product's introduction.

*Figure A8:* Distribution of imitation lag (days between third-party product introduction and private-label entry) by seller revenue quintile

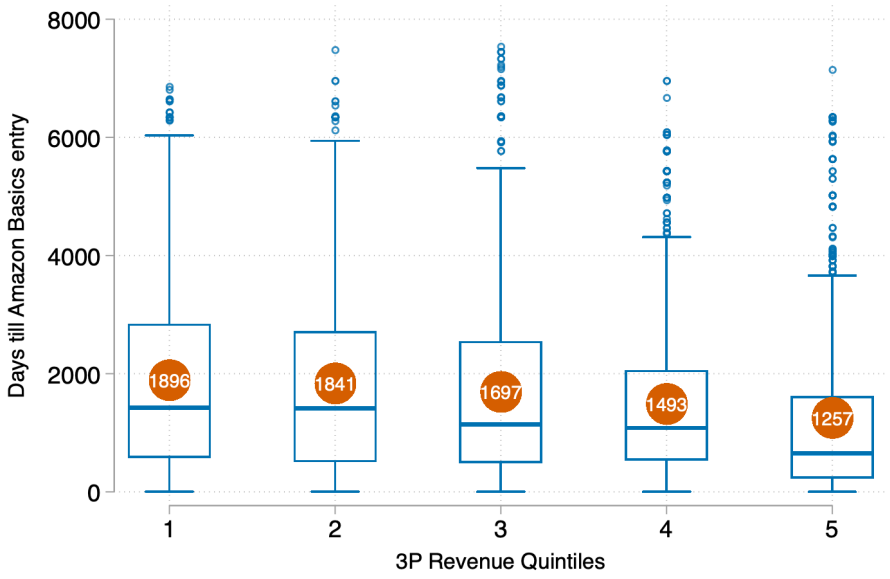
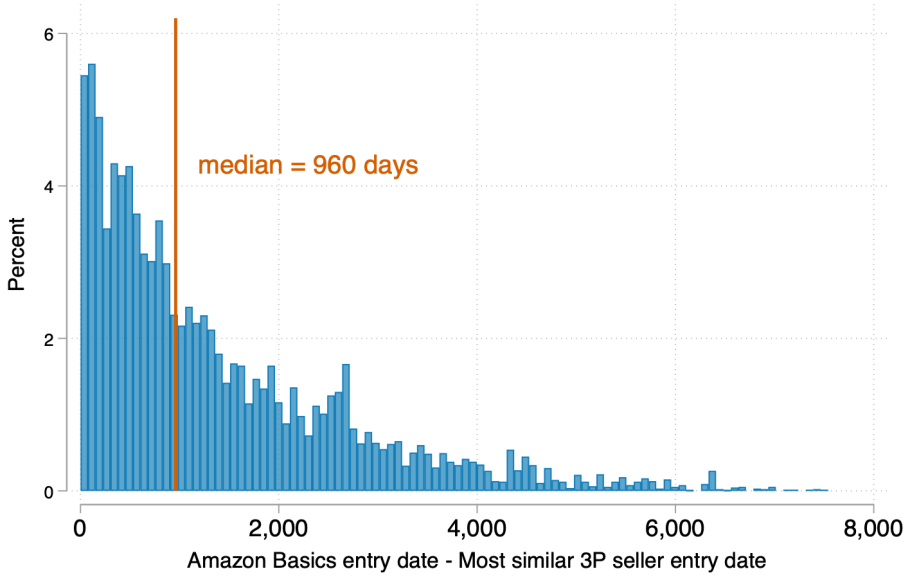



*Figure A9:* Distribution of days between third-party product introduction and similar private-label entry



# E Robustness of Seller Count Decline and Entry-Exit Decomposition

## E.1 Sensitivity to Category Size

The baseline estimates in Table 7 include all product categories, many of which are small or near-empty at baseline. To assess whether the seller count results are driven by these marginal categories, we re-estimate the Callaway and Sant'Anna (2021) model restricting to categories with at least 5, 20, or 50 products at baseline (measured as the pre-treatment mean). Table A5 reports the results for log total sellers.

The U.S. estimate attenuates substantially—from $-0.90$ to $-0.47$—as small categories are excluded, indicating that marginal categories with few baseline products contribute disproportionately to the U.S. result. At the most restrictive cutoff ($\geq 50$ products), the U.S. ATT of $-0.47$ implies that treated categories have 37% fewer products than they would have absent Amazon's entry, which remains economically meaningful and highly significant. By contrast, the estimates for Germany, the U.K., and Canada are stable or even increase with stricter cutoffs, suggesting that the competitive effects in these markets are robust across the category size distribution.

## E.2 Entry-Exit Decomposition: Raw Counts

The decline in total products combines two margins: increased exit and reduced entry. Table A6 reports the Callaway and Sant'Anna (2021) estimates for raw product entries, exits, and net change.

The net change is significantly negative in all four markets. However, neither entries nor exits are individually significant in the U.S. or Germany, while exits are significant in the U.K. and Canada. The net effect is more precisely estimated than its components because entries and exits are highly positively correlated:

***Table A5:*** Sensitivity of log total sellers ATT to minimum category size

| | Amazon Marketplace | | | |
|---|---|---|---|---|
| Minimum Baseline Size | U.S. | DE | U.K. | CA |
| All categories (baseline) | $-0.901^{***}$ | $-1.650^{***}$ | $-0.990^{***}$ | $-0.974^{***}$ |
| | (0.075) | (0.084) | (0.100) | (0.084) |
| | [8,351 categ.] | [5,561 categ.] | [5,874 categ.] | [6,382 categ.] |
| | [$N = 317{,}281$] | [$N = 208{,}470$] | [$N = 222{,}513$] | [$N = 242{,}238$] |
| $\geq 5$ products | $-0.697^{***}$ | $-1.394^{***}$ | $-1.129^{***}$ | $-0.920^{***}$ |
| | (0.057) | (0.077) | (0.076) | (0.081) |
| | [3,842 categ.] | [3,449 categ.] | [2,155 categ.] | [2,408 categ.] |
| | [$N = 256{,}887$] | [$N = 162{,}141$] | [$N = 134{,}608$] | [$N = 158{,}792$] |
| $\geq 20$ products | $-0.572^{***}$ | $-1.350^{***}$ | $-1.268^{***}$ | $-1.041^{***}$ |
| | (0.058) | (0.079) | (0.083) | (0.078) |
| | [2,912 categ.] | [2,503 categ.] | [1,390 categ.] | [1,572 categ.] |
| | [$N = 219{,}606$] | [$N = 137{,}589$] | [$N = 99{,}774$] | [$N = 118{,}515$] |
| $\geq 50$ products | $-0.469^{***}$ | $-1.317^{***}$ | $-1.316^{***}$ | $-1.076^{***}$ |
| | (0.062) | (0.084) | (0.089) | (0.079) |
| | [2,304 categ.] | [1,887 categ.] | [908 categ.] | [1,063 categ.] |
| | [$N = 181{,}406$] | [$N = 110{,}573$] | [$N = 69{,}879$] | [$N = 85{,}090$] |

*Notes:* Each cell reports the aggregate *ATT* from the Callaway and Sant'Anna (2021) estimator on log(total sellers). Baseline size is the pre-treatment mean number of products in the category. Observation counts ($N$) in brackets show how the sample shrinks as the minimum size cutoff increases. The U.S. starts with 8,351 categories and drops to 2,304 at the $\geq 50$ cutoff; other markets follow a similar pattern. The U.S. estimate attenuates from $-0.90$ to $-0.47$ as small categories are excluded, while estimates for Germany, the U.K., and Canada are stable or increase. All estimates significant at the 1% level. Bootstrap standard errors (200 iterations) in parentheses.

***Table A6:*** Decomposition of product count changes: entry, exit, and net change

| | Amazon Marketplace | | | |
|---|---|---|---|---|
| Outcome | U.S. | DE | U.K. | CA |
| Entries | 122.0 | −1,026.9 | 23.9 | 94.5 |
| | (429.5) | (960.3) | (82.4) | (198.6) |
| Exits | 683.2 | 233.7 | 314.6*** | 508.4** |
| | (428.8) | (503.9) | (63.3) | (205.9) |
| Net Change | −561.2** | −1,260.6** | −290.7*** | −413.9* |
| | (219.6) | (590.4) | (80.9) | (229.4) |
| Observations | 370,348 | 233,270 | 222,513 | 242,238 |
| Categories | 4,140 | 3,656 | 2,493 | 2,610 |

*Notes:* Each cell reports the aggregate *ATT* from the Callaway and Sant'Anna (2021) estimator. Entries and exits are raw counts of products entering or leaving the category per month. Net change equals entries minus exits. Observation counts differ from Table 7 because the extensive margin panel (tracking entry and exit flows) covers a broader set of category-months than the intensive margin panel (which requires non-missing seller and price data). The net effect is precisely estimated despite noisy components because entries and exits are positively correlated within categories. Standard errors (bootstrap, 200 iterations) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

$\text{Var}(\text{Entries} - \text{Exits}) = \text{Var}(\text{Entries}) + \text{Var}(\text{Exits}) - 2\,\text{Cov}(\text{Entries}, \text{Exits})$, and the large positive covariance substantially reduces the variance of the difference. The event studies for individual components exhibit noisy pre-trends, so we interpret these results as supplementary.

## E.3 Entry-Exit Decomposition: Rates Normalized by Baseline

Raw entry and exit counts are noisy because they scale with category size. To address this, we normalize by baseline category size: for each category, we divide monthly entries and exits by the pre-treatment mean number of active products, yielding entry rates and exit rates. Table A7 reports the results.

The rate-based decomposition reveals a striking pattern: both entry and exit rates increase significantly after private-label entry in most markets. In the U.S., the entry rate rises by 0.24 and the exit rate rises by 0.34 (both significant at the 1% level). The exit rate

***Table A7:*** Entry and exit rates (normalized by baseline category size)

| | Amazon Marketplace | | | |
|---|---|---|---|---|
| Outcome | U.S. | DE | U.K. | CA |
| Entry Rate | 0.243*** | 0.418 | 0.997*** | 1.490*** |
| | (0.043) | (0.261) | (0.258) | (0.262) |
| Exit Rate | 0.339*** | 0.762*** | 1.108*** | 1.654*** |
| | (0.031) | (0.266) | (0.255) | (0.265) |
| Net Change Rate | −0.096*** | −0.345*** | −0.111*** | −0.164*** |
| | (0.027) | (0.073) | (0.041) | (0.051) |
| Observations | 276,100 | 178,001 | 160,833 | 180,582 |
| Categories | 4,090 | 3,629 | 2,470 | 2,589 |

*Notes:* Entry rate and exit rate are defined as the number of products entering (exiting) a category in a given month divided by the pre-treatment mean number of active products in that category. Categories with fewer than 5 baseline active products are excluded, reducing the sample relative to Table A6. Net change rate equals entry rate minus exit rate. Standard errors (bootstrap, 200 iterations) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

increase exceeds the entry rate increase, producing a significantly negative net change rate of $-0.10$. This pattern is consistent across all four markets, with the net change rate ranging from $-0.10$ (U.S.) to $-0.35$ (Germany).

These results indicate that private-label entry does not simply suppress entry or accelerate exit in isolation. Rather, it increases competitive churn—more products enter and more products exit treated categories—but exits outpace entries, generating the net decline observed in the intensive-margin estimates. This turbulence is consistent with Amazon's private-label entry triggering a wave of competitive repositioning: some third-party sellers exit, while others enter (perhaps at different price points or with differentiated products), but the overall effect is a net reduction in the product stock.