Writing with Generative AI: Consumer Demand in Online Marketplaces

Nan Chen Hsin-Tien Tsai*

November, 2026

Abstract

This paper examines how generative AI affects market competition through its use in product-listing text. We study Amazon's rollout of seller-facing generative-AI listing tools in September 2023, combining data on product entry, consumer clicks and purchases, and historical listing text. Our identification exploits an institutional feature of Amazon: sellers of single-listing products have greater ability to revise and improve the product page. We therefore use markets' pre-launch shares of single-listing products as a measure of treatment intensity. Markets with higher treatment intensity experience both more product entry and greater concentration of clicks and purchases among the top products. The concentration effect is smaller in categories where sellers have limited ability to rewrite listing text and larger for descriptive searches. In more exposed markets, seller reputation and review histories become more important for demand, while lower prices matter more for foreign products, suggesting that generative AI shifts competition toward non-text attributes. Consistent with this reallocation, the increase in concentration is accompanied by lower prices and slightly higher ratings rather than worse observable outcomes for consumers.

Keywords: Generative AI; Online Marketplaces; Product Entry; Market Concentration; Product Information; Search and Discovery

*Chen: Department of Information Systems and Analytics, National University of Singapore, nan@nus.edu.sg. Tsai: Department of Economics, National University of Singapore, ttsai@nus.edu.sg. This paper receives financial support from a tier-1 grant provided by the Ministry of Education in Singapore.

1 Introduction

Generative AI has made polished written content nearly free to produce. One of its first large-scale commercial uses is writing the product listings through which sellers compete in online marketplaces. Cheaper listing text lowers the cost of introducing a product and may therefore encourage entry. It also changes what listings tell consumers. When every seller can produce fluent and complete text, text may stop distinguishing products, and competition may shift toward attributes that remain costly to acquire. The net effect on market structure is an open empirical question: more products may compete for consumers' attention, yet attention may end up more concentrated.

This paper studies these questions using Amazon's rollout of seller-facing generative-AI listing tools, which began in September 2023. The tools draft a product's title, description, and structured attributes from a short prompt, and sellers accept the output with little or no editing most of the time.[1] We ask how the tools affect product entry, the concentration of consumer demand among top products, and the attributes that determine which products capture that demand.

We combine four data sources. Amazon Brand Analytics reports weekly click and conversion shares of the three most-clicked products for roughly one million search terms, providing a direct measure of product-level concentration within each search query. Entry data that we collect track more than three million products in the Home & Kitchen and Office categories, including entrants that never become prominent enough to appear in Brand Analytics. Historical product-page snapshots from the Wayback Machine let us observe within-product text revisions for more than 94,000 products from before the introduction of ChatGPT through 2025. Keepa provides prices, sales ranks, ratings, review counts, and seller location.

The central identification challenge is that sellers' use of the tool is unobserved and endogenous. We instead exploit an institutional feature of Amazon. Every product has a single detail page shared by all sellers that offer it, and Amazon, not the sellers, decides whose submitted text is displayed. When a product is offered by many sellers, no individual seller controls the page, and sellers compete

[1] Amazon reported more than 100,000 users by March 2024, more than 400,000 by September 2024, and more than 900,000 by May 2025, and that sellers accepted AI-generated content with little or no editing about 90 percent of the time. Section 2 documents the rollout.

through price and fulfillment rather than text. When a product has a single seller, that seller is the main contributor and can revise the page using the tool. We therefore measure a market's exposure to the tool by its pre-launch share of single-listing products, defined at the keyword level for demand outcomes and at the leaf-category level for entry. Because this measure is predetermined, it does not depend on sellers' subsequent adoption choices. Identification requires markets with different treatment intensity to have followed parallel outcome trends absent the rollout. We estimate continuous-treatment difference-in-differences specifications and report event studies throughout. Because Amazon later introduced consumer-facing AI tools, beginning with the Rufus shopping assistant in February 2024, we also report estimates using a pre-Rufus window ending in January 2024 for monthly outcomes and in the fourth quarter of 2023 for quarterly outcomes.

We first show that the exposure measure captures use of the tool. We detect generative-AI text using words that become disproportionately common in revisions of Amazon listings after the introduction of ChatGPT, together with em and en dashes; results are similar using words that prior work identifies as overused by ChatGPT. After the launch, products with fewer listings revise their text more often and add more generative-AI markers, while the corresponding coefficients are close to zero before the launch. The same pattern holds across markets: the post-launch increase in the share of clicked products carrying generative-AI markers rises with the keyword's pre-launch single-listing share.

Our first main result is that entry rises sharply in more exposed markets. Before the release of Rufus, a one-standard-deviation increase in a category's single-listing share raises entry by 2.8 products per quarter, 55 percent of the pre-launch mean, and the effect continues to grow through 2025. The additional entrants are almost entirely single-listing products, precisely the products for which Amazon's rules give sellers control over the page. Most are sold by sellers located in China. Incumbents' sales ranks worsen in the same markets, consistent with more intense competition.

Our second result is that, despite this entry, demand becomes more concentrated. A one-standard-deviation increase in exposure raises the combined click share of the top three products by 1.4 percent and their combined conversion share by 6.5 percent, and most of each effect is already present before Rufus. Two patterns tie the effect to listing text. It is much smaller in media categories,

where product pages are built from catalog metadata that sellers cannot rewrite. And it increases with how descriptive the search query is. The effect is small for brand-name searches, in which consumers have already chosen a brand, and largest for queries that specify three or more product attributes, where listing text has the most scope to determine which product matches.

Our third result is that concentration does not entrench pre-launch leaders. The portion of the top-three share held by products that were at the top before the launch falls, while the portion held by other products rises, and the increase in concentration arises in keywords that had no dominant product before the launch. The tool thus weakens the persistence of pre-existing advantages while concentrating demand among a new set of products that rise to the top.

We then ask why cheaper text concentrates demand. In more exposed markets, listings adopt more generative-AI language, become less repetitive, and move toward the text of the keyword's pre-launch top product; these changes are concentrated among domestic sellers. As text converges, it becomes a weaker basis for differentiation: products whose listings become more similar to their competitors' lose clicks and conversions. At the same time, non-text attributes gain importance. Products sold by highly rated sellers perform better; among domestic sellers, highly rated products gain, while among foreign sellers they lose. Products with few reviews and older products also lose, and lower prices become more important for foreign sellers. The most-clicked products have average prices about 2 percent lower and slightly higher average ratings in more exposed markets, suggesting that the increase in concentration is not accompanied by worse observable consumer outcomes.

Taken together, the results show that a technology that makes one input to competition nearly free can expand product entry while concentrating demand. The two effects are two sides of the same change: when listing text no longer distinguishes products, demand flows to the products with the strongest ratings, reputations, and prices, and those products need not be the ones that led the market before.

Related Literature. This paper contributes to research on the market-level effects of generative AI. Our closest comparison is Goldberg and Lam (2026), who study how AI-generated creative goods affect entry, product supply, and demand in an

online marketplace. Their setting highlights market expansion and the displacement of existing producers when AI can generate the traded good itself. Related work examines creative production, online labor demand, and entrepreneurship (Kim et al., 2026; Teutloff et al., 2025; Hui et al., 2024; Cai et al., 2025), while Fang et al. (2025) study generative AI in online retail. We also build on research on firm-level AI adoption, productivity, growth, and industry structure (Czarnitzki et al., 2023; Babina et al., 2024; McElheran et al., 2024; Lu et al., 2024) and on productivity gains in writing, customer service, and software development (Noy and Zhang, 2023; Brynjolfsson et al., 2025; Peng et al., 2023; Cui et al., 2026). We examine AI-assisted production of the descriptions through which products compete, tracing how cheaper content production changes product entry and the allocation of consumer attention and purchases across competing listings.

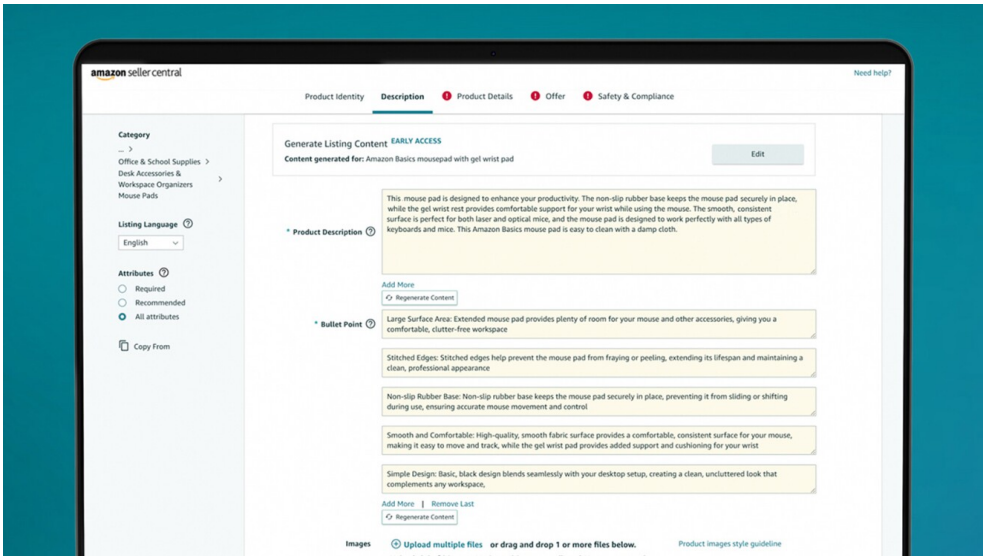
The paper also contributes to research on entry, variety, and competition in digital markets. Research on digitization shows how lower production costs can expand product variety and generate consumer benefits (Aguiar and Waldfogel, 2018; Waldfogel, 2017), while lower experimentation costs can change the financing of entrepreneurial entry (Ewens et al., 2018). Platform-provided analytics can also improve seller outcomes (Bar-Gill et al., 2024). Models of search and product differentiation emphasize that market outcomes depend on how consumers discover and compare products (Anderson and Renault, 1999; Bar-Isaac et al., 2012). We bring entry and demand allocation together: markets more exposed to the AI tool experience greater product entry alongside higher concentration of clicks and conversions, so an expansion in the number of available products need not disperse consumer attention or purchases.

Finally, our findings relate to information and signaling in markets. When producing effective written communication is costly, its quality may convey information about the sender (Spence, 1973). Evidence on algorithmic writing assistance shows that changes in presentation can affect market outcomes (Wiles et al., 2025), and recent studies examine how generative AI changes the informativeness of written signals and the use of alternative quality indicators (Galdin and Silbert, 2025; Cui et al., 2025). Experimental evidence also shows that AI assistance can make independently produced outputs more similar (Doshi and Hauser, 2024). We study these forces in product-market competition. Listing text converges toward that of pre-launch top products, while ratings, accumulated reviews, and prices be-

come more predictive of top positions, consistent with a change in the dimensions of differentiation when polished product descriptions become easy to produce.

The paper proceeds as follows. Section 2 describes Amazon's generative-AI listing tools, and Section 3 presents the data and our measure of generative-AI text. Section 4 develops the identification strategy and shows that ownership structure predicts use of the tool. Section 5 presents the results on product entry and incumbent sales, and Section 6 the results on market concentration, including the heterogeneity by rewriting ability and query language and the analysis of pre-launch dominance. Section 7 examines how listing text changes, which attributes gain importance, and the ratings and prices of top products. Section 8 concludes.

2 Background



Note: The figure illustrates Amazon's generative-AI listing tool. The interface allows sellers to generate product-listing content from a short product description, such as a mousepad with a gel wrist pad. The tool can generate listing elements including the product title, product description, and structured product details, which sellers can review before submitting.

Figure 1: Example of Amazon's seller-facing generative-AI listing tool.

2.1 Amazon's Generative-AI Listing Tools

Amazon began rolling out seller-facing generative-AI listing tools on September 13, 2023.[2] Within Seller Central, sellers could enter a short product description, for example, "Amazon Basics mousepad with gel wrist pad", and receive a draft title, product description, and structured product details. Sellers could then edit, refine, or submit the output. By generating catalog attributes as well as prose, the tool shifts listing creation from de novo writing and field-by-field data entry toward verification and exception correction. Figure 1 illustrates this interface. The resulting workflow is expected to reduce the fixed labor required to introduce a product while potentially producing more complete information for search indexing and customer discovery.

Amazon subsequently expanded both the accepted inputs and the range of listing tasks. By March 2024, sellers could generate content from a product image, and Amazon began rolling out a workflow that converted an existing direct-to-consumer product-page URL into an Amazon-ready listing.[3] Images could be used to infer details such as color, keywords, and other attributes. Later expansions added bulk generation from sparse spreadsheets and recommendations for improving existing listings. Thus, short text, images, and URLs can generate individual listings; spreadsheets can generate multiple listings; and existing listings can be evaluated for missing or improvable content.[4]

Appendix A describes how Amazon's use of generative AI progressed from product-listing creation to a broader infrastructure for seller productivity, content optimization, and customer discovery. Related tools include Project Amelia, Enhance My Listing, product-image and video generators, and the Rufus AI shopping assistant.

[2]Amazon, "Amazon launches generative AI to help sellers write product descriptions," September 13, 2023. https://www.aboutamazon.com/news/small-business/amazon-sellers-generative-ai-tool

[3]Amazon, "Amazon selling partners can now access even more generative AI features to create high-quality product listings," March 13, 2024. https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-powered-product-listings

[4]See https://sell.amazon.com/blog/amazon-listing-ai and https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-seller-growth-shopping-experience.

2.2 Generative-AI Listing Tool Adoption

Amazon's own disclosures point to rapid growth in the use of its generative-AI listing tools. The company reported more than 100,000 users in March 2024, more than 400,000 globally in September 2024, and more than 900,000 by May 2025 who had used the tools at least once.[5] For context, using two million active sellers as a rough benchmark, these reported user counts would correspond to approximately 5%, 20%, and 45% of the active-seller base, respectively.[6]

It is important to note that these statistics cover users of Amazon's own tools embedded in its listing workflow. They do not capture sellers who use specialized third-party software or general-purpose language models outside Seller Central. Sellers who rely only on external tools would therefore not be counted, so Amazon's reported figures can be viewed as a lower bound on overall generative-AI adoption. Consistent with this, a survey of nearly 2,000 Amazon businesses across more than 20 marketplaces conducted around January 2024 found that 34% of Amazon sellers and brands used AI to write or optimize listings, which is higher than the share implied by Amazon's internal-tool count at that time.[7]

Generative-AI Listing Quality. Amazon reports the extent to which sellers accept AI-generated content without substantial revision. In March 2024, sellers accepted suggested attributes nearly 80% of the time with minimal edits; by May 2025, Amazon reported that sellers accepted AI-generated content with little or no editing approximately 90% of the time.[8] These statistics measure how much editing sellers make after using the tool and suggest that sellers often find the

[5]For the March 2024 disclosure, see https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-powered-product-listings. For the September 2024 and May 2025 figures, see https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-seller-growth-shopping-experience.

[6]Third-party estimates suggest that Amazon had roughly two million active sellers during this period. For example, Marketplace Pulse estimates that the global active-seller population declined from 2.4 million in 2021 to about 1.65 million by the end of 2025; see https://www.marketplacepulse.com/articles/amazon-seller-registrations-hit-decade-low-in-2025. Other estimates are closer to 1.8 million; see https://capitaloneshopping.com/research/amazon-marketplace-statistics/.

[7]See Jungle Scout, https://www.junglescout.com/resources/reports/amazon-seller-report-2024/.

[8]See https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-powered-product-listings and https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-seller-growth-shopping-experience.

generated content usable with little revision.

Amazon's later seller guidance also states that the tools can generate more than 70% of the product attributes required for a listing. This suggests that AI can improve listing completeness by filling structured fields that sellers might otherwise omit, which may support search indexing, product filtering, and customer discovery.[9] Amazon also reports a 40% increase in its internal measure of overall listing quality when sellers use the tools.[10] This measure captures more than writing quality alone: Amazon describes it as a broader data-quality measure based on seller-submitted listings, manual audits, seller feedback, and listing-acceptance rates.[11]

3 Data

The data used in the analysis come primarily from four sources. First, the Brand Analytics Data provide information on the performance of top products for approximately one million of the most popular search terms on Amazon. Second, we manually collect data on new product entry and existing incumbents in the Home & Kitchen category. Third, we obtain product descriptions, sales ranks, and prices from Keepa. Finally, we use the Wayback Machine to track historical changes in product detail pages over time.

Brand Analytics Data. Amazon Brand Analytics reports the three top products, which we refer to as the top-three products, for each of Amazon's top search terms in each week. For each product, we observe the product ID, the title displayed that week, weekly click share, and weekly conversion share. A product's click share is the number of clicks it receives from searches for a given term divided by the total number of clicks generated by that search term in that week; its conversion share is defined analogously using purchases attributed to that term.

[9] Amazon, "How to use Amazon AI to create product listings," November 3, 2025. https://sell.amazon.com/blog/amazon-listing-ai

[10] See https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-seller-growth-shopping-experience.

[11] See the discussion of AI-assisted listing quality in https://sell.amazon.com/blog/amazon-listing-ai.

Panel A of Table 1 reports summary statistics for the Brand Analytics Data. The average top-three product captures 13.6% of its search term's clicks, implying that the top three products jointly account for roughly 40% of all clicks and about 30% of all conversions for the average search term. Our sample runs from March 25, 2023 through the end of 2025.[12] The median product appearing in a top-three observation has approximately 1,760 reviews and a 4.5-star rating, while the mean review count approaches 10,000. Most products are well established: the median product is nearly five years old. We measure a product's number of listings as its average daily number of new-condition listings during the pre-launch period, April 1 to September 9, 2023, and classify it as single-listing if this average is at most one. About 39% of observations are single-listing products. The number of listings is highly skewed: the median product has 1.3 listings, while the mean is 5.4 because some products are sold by many sellers. About 6.6% of observations are media products, such as books, movies, recorded music, and video games.[13]

Entry Data. We focus on categories with relatively high rates of new product introduction, including Kitchen & Dining, Home Decor, and Office Products. In each iteration, we systematically use Keepa.com to identify the product IDs of recently added products in these focal categories. Newly identified products are then added to the set of products that we subsequently monitor; see Appendix B for details of the collection procedure. A key advantage of these Entry Data relative to the Brand Analytics Data is that they capture new products regardless of their subsequent success. In particular, the data include entrants that never become sufficiently prominent to appear among the top-three products for any search term. Indeed, only about 4.4% of products in the Entry Data ever appear among the top-three products for any search term in the Brand Analytics Data. The Entry Data therefore allow us to study the broader distribution of product introductions rather than only the relatively successful products observed in Brand Analytics.

After identifying the entrant product IDs, we separately collect detailed product

[12]The raw data extend back to 2021. However, Amazon restructured the Brand Analytics reports on March 25, 2023, creating substantial level breaks in several outcome variables, particularly the share measures. We therefore restrict the analysis to the period from March 25, 2023 onward.

[13]The prevalence of single-listing products is very similar among media and non-media keywords: the average single-listing share is 42 percent for media keywords and 44 percent for non-media keywords.

information for both entrants and incumbents. From May 2023 through November 2024, we scrape and track product-level information for these products over time. The resulting panel covers approximately 3.1 million products and includes products sold by more than 160,000 independent third-party sellers.

Panel B of Table 1 reports summary statistics for the Entry Data. Products in the Entry Data are substantially newer and smaller than those in the Brand Analytics Data: the median product has only 62 reviews, and ratings are observed for approximately one-quarter of products, as many newly introduced products have not yet received a rating. The distribution of time since introduction reflects the coexistence of two distinct populations: incumbent products have been listed for a median of approximately fifteen years, whereas entrants have been listed for roughly seven months. The Entry Data contain no media categories. Of approximately 3.12 million products, 1.64 million were listed before the September 2023 launch of Amazon's generative-AI listing tools and 1.47 million (47.3%) entered afterward.

Because entrants did not exist before the launch, we measure their number of listings over the first 180 days after listing. Entrants almost always arrive as single-listing products: 94% are single-listing, compared with 39% in the Brand Analytics Data, which consists of larger and more established products. The average number of listings among entrants is 1.05.

Auxiliary data. We supplement these data with product characteristics and seller information from Keepa.com. Keepa provides high-frequency time series of prices, sales ranks, ratings, and review counts, as well as product-level information such as listing date, title, product description, images, and brand. Sales rank refers to Amazon's Best Sellers Rank, an ordinal ranking of products within a category based on recent sales and updated at high frequency. The best-selling product is ranked first, so a lower sales rank indicates better sales performance.[14]

We also collect historical snapshots of Amazon product pages from the Wayback Machine to track changes in listing content over time. The collection covers 94,317 products from the Brand Analytics Data, including 12,572 media products, and

[14]Because sales quantities are approximately log-linearly related to sales rank, we use the log of sales rank as our measure of sales performance. Changes in log sales rank can therefore be interpreted as an approximate negative proportional change in sales volume.

contains 3.5 million parsed captures from January 2020 through August 2026. We extract the product title and description from each capture. The median product has 9 captures, and 67% of products have at least one capture both before and after the September 2023 launch of Amazon's generative-AI listing tool. For the text-revision analyses, we compare each product's last capture in a quarter with its most recent preceding capture, yielding 452,470 product-quarter comparisons across 80,693 products from 2021 through 2025. See Appendix C for the Wayback Machine Panel summary.

3.1 Measuring Generative-AI Text

We use two proxies for generative-AI-associated writing, based on punctuation and word choice. These measures capture changes in listing language but do not identify whether an individual listing was produced using Amazon's tool. The first is based on punctuation. We flag listings containing an em dash or en dash, which large language models use substantially more often than human sellers. Dash use therefore provides a simple proxy for AI-associated writing.

The second measure is based on words that appear more often in AI-written listings; see Appendix D for the list and their frequencies. Rather than choosing these words ourselves, we identify them from the data. Using revisions of the same non-media product in the Wayback Machine Panel, we compare words added after the release of ChatGPT with those added before.[15] We retain words whose share among added words is at least three times higher after the release of ChatGPT and that are added at least 100 times after the release.[16]

[15] We exclude media products, such as books, movies, recorded music, and video games, because their product pages are largely constructed from catalog metadata that sellers have limited ability to modify (see Appendix F).

[16] For each pair of consecutive Wayback Machine captures of the same non-media product whose title or product description changed, we record the words added in the later capture. For each word, we calculate its share of all added words in revisions made before and after the release of ChatGPT. To ensure that the selected words are not driven by a small number of products, we randomly split products into two halves and require the threefold increase in both halves. We use the Wayback Machine Panel because it records revisions of the same listing, so the selected words reflect changes in how sellers write rather than cross-product differences. The procedure uses the release of ChatGPT rather than the release of Amazon's generative-AI listing tool and does not use any of the outcomes studied later. Because the words are selected using the timing of ChatGPT, Appendix G shows similar results using words that prior studies document as overused by ChatGPT, while Section D.1 shows similar results using words selected based on dash use rather than timing.

Variable	Mean	SD	Median
Panel A: Brand Analytics Data			
Click share	0.136	0.112	0.103
Conversion share	0.098	0.135	0.059
Title length (chars)	123.8	55.0	130
Rating (stars)	4.48	0.29	4.50
Review count	9,793	29,434	1,762
ln(sales rank)	8.63	2.36	8.82
Price ($)	51.1	363.4	20.0
Time since introduction (months)	75.9	62.7	56.0
Single-listing	0.394	0.489	0
Number of listings	5.43	11.78	1.31
Media	0.066	0.248	0
No. of Observations		68,820,359	
Panel B: Entry Data			
Title length (chars)	142.9	42.9	151
Rating (stars)	4.41	0.60	4.50
Review count	1,286	6,329	62
ln(sales rank)	2.92	1.65	2.94
Price ($)	39.1	188.0	18.3
Time since introduction (months)	93.6	79.9	64.7
Single-listing (entrants)	0.937	0.244	1
Number of listings (entrants)	1.05	0.81	1.00
Entrant	0.473	0.499	0
No. of Observations		3,115,165	

Notes: The table reports summary statistics for the two datasets used in the analysis. Panel A summarizes the Brand Analytics Data in 2024 at the weekly product level. Single-listing products are defined as those with at most one listing, where the number of listings is measured as the average daily number of new-condition listings during the pre-launch period. Panel B summarizes the Entry Data at the product level. Entrants are defined as products first listed after the September 13, 2023 launch of Amazon's generative-AI listing tools. In Panel B, the number of listings for entrants is measured as the average daily number of new-condition listings over the first 180 days after listing. Product ratings, review counts, and sales ranks are matched from Keepa data.

Table 1: Summary Statistics

Table A3 in Appendix D lists all 49 words that meet these criteria. We exclude nine words that do not describe writing style: words tied to particular products, places, or times of day, such as *pfas*, *feline*, *apartments*, and *mornings*; common everyday words, such as *shared* and *differ*; and the German word *und*. This leaves 40 words, including *ensuring*, *crafted*, *effortless*, *elevate*, and *meticulously*.

We classify a listing as containing generative-AI markers if its title or product description contains at least one of these words, an em dash, or an en dash. Overall, approximately 49% of listings in the Entry Data and 32% of listings in the Brand Analytics Data contain at least one generative-AI marker.

AI-Generated Text Adoption. In Figure 2, we examine generative-AI text adoption both within existing products and across entry cohorts. We classify a listing as generative-AI-marked if its text contains an em-dash, an en-dash, or any word from our 40-word AI vocabulary.

For products in the Brand Analytics Data, we use the Wayback Machine Panel, which records the same listing's text at successive points in time. We measure generative-AI marker use as the number of occurrences of the 40 words in Table A3 per 1,000 words of the listing's title and product description, and estimate

$$y_{it} = \alpha_{i,d_{it}} + \sum_{q \neq q_0} \gamma_q \, \mathbf{1}[t = q] + \varepsilon_{it}, \tag{1}$$

where y_{it} is the outcome for product i in quarter t, measured using its last Wayback Machine capture in that quarter, and d_{it} indicates whether that capture includes the full product description. Because the Wayback Machine does not always capture the full product description, the product fixed effects $\alpha_{i,d_{it}}$ are allowed to differ by d_{it}. The coefficients γ_q therefore capture within-product changes over time while accounting for systematic differences associated with whether the full product description is captured. They are measured relative to q_0, the third quarter of 2022, the last quarter before the introduction of ChatGPT. Standard errors are clustered at the product level.

Figure 2a presents the series. Before the introduction of ChatGPT, the series is flat, with all estimates within 0.04 markers per 1,000 words of the baseline. Generative-AI marker use begins to rise in the second quarter of 2023 and is 0.12 markers per 1,000 words above the baseline in the quarter of Amazon's generative-AI listing tool release, consistent with some adoption following the introduction of ChatGPT. The increase then accelerates after the release of the generative-AI listing tool: Generative-AI marker use is 0.29 markers per 1,000 words higher by the end of 2023, 0.74 higher by the end of 2024, and 1.15 higher by the end of 2025, relative to a baseline of 0.8 markers per 1,000 words in the third quarter of 2022.

These within-product increases indicate that generative-AI markers increasingly appear in revisions to existing listings.

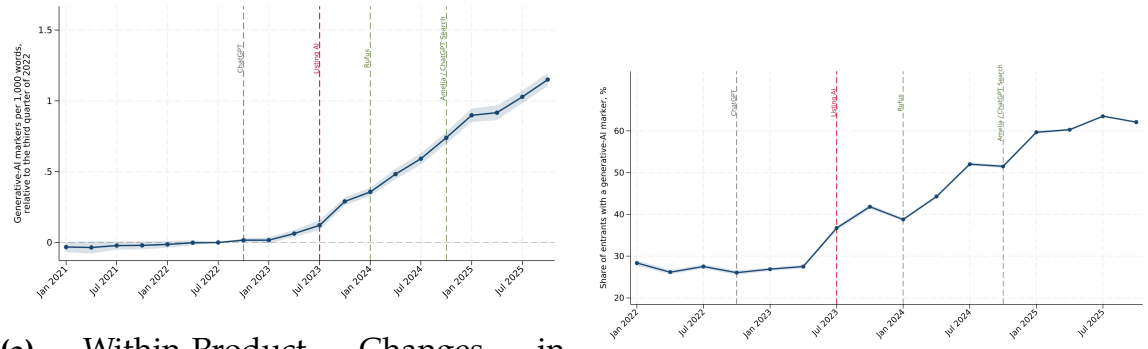
For entrants, product descriptions for the full sample were collected from Keepa in 2026, so we do not observe within-product text revisions following the introduction of generative-AI tools. Our main variation in text characteristics is therefore cross-sectional. We compare products introduced before and after these tools became available, since the text of newer listings is more likely to reflect the generative-AI technology available at the time of their creation.

Figure 2b plots the share of products whose listings contain a generative-AI marker by quarter of introduction. We use quarterly cohorts to ensure a sufficient number of observations in each cohort. The share of products with generative-AI markers increases sharply across cohorts introduced after these tools became available, more than doubling from about 28% to nearly 60% for products introduced in 2025. These cohort patterns complement the within-product evidence of increasing generative-AI text adoption.

Robustness to Marker Construction. Section D.1 shows that our results do not depend on how the generative-AI markers are selected. Within products in the Wayback Machine Panel, generative-AI marker use also increases after the release of the generative-AI listing tool when the words are selected based on dash use rather than timing, either using the 70 words that are at least twice as common in Entry Data listings with dashes as in listings without dashes, or selecting words from one half of the Wayback Machine products and applying them to the other half.

4 Empirical Strategy

Our identification relies on two sources of variation: the rollout of Amazon's generative-AI listing tool and an institutional feature of Amazon that determines which sellers can use the tool to change what consumers see. This section describes that feature, defines the resulting market-level measure of treatment intensity and the specifications we estimate, and presents evidence that the measure captures differential use of the tool.



(a) Within-Product Changes in Generative-AI Marker Use in the Brand Analytics Data

(b) Share of New Products with Generative-AI Markers in the Entry Data

Notes: The figure documents evidence of generative-AI text adoption. Generative-AI markers are defined in Section 3.1. Figure 2a uses Wayback Machine captures of products in the Brand Analytics Data and plots the coefficients γ_q from Equation (1), using the number of generative-AI marker occurrences per 1,000 words of the title and product description as the outcome. The resulting series is normalized to zero in the third quarter of 2022, the last quarter before the introduction of ChatGPT. Figure 2b uses the Entry Data and plots the share of products classified as generative-AI-marked by quarterly cohort, defined by each product's first listing date. A product is classified as generative-AI-marked if its title or product description contains an em-dash, an en-dash, or any word from the 40-word generative-AI marker list described in Section 3.1. Bands show 95% confidence intervals; standard errors in Figure 2a are clustered at the product level. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure 2: Adoption of Generative-AI Markers

4.1 Editorial Control on Shared Product Pages

Each product on Amazon is identified by a unique product code (ASIN) and displayed on a single detail page. Sellers list offers on this page but do not own it. A seller offering an existing product can attach an offer to its detail page, acting as a reseller. All offers associated with that product share the same page, URL, title, product description, images, and other structured attributes. Amazon's policy explicitly states that a detail page "is not owned by any one seller" and prohibits duplicate pages for the same product.[17] Thus, sellers share a common product page but compete for sales through their individual offers, with many purchases going to the offer featured in the Buy Box. Competition among sellers of the same product therefore centers on price, fulfillment, and Buy Box eligibility rather than listing text.

[17]See Amazon Seller Central, "Product detail page rules" and "Duplicate product detail pages," https://sellercentral.amazon.com/help/hub/reference/external/G200390640.

Because multiple sellers may submit values for the same product attribute, Amazon treats these submissions as *contributions* and determines which information appears on the shared detail page. Amazon's public guidance indicates that contributions from brand owners enrolled in Brand Registry generally receive priority, and that a Brand Representative's contributions take precedence over those of a Reseller when both have assigned brand selling roles. Amazon also considers seller performance, including sales volume, refund rates, buyer feedback, and A-to-z Guarantee claims, when evaluating contributions. The ability to submit listing text therefore does not guarantee control over the text displayed to buyers: a seller's proposed changes may not appear if Amazon selects another contributor's information. When changes cannot be made through self-service tools, sellers may seek support review by providing evidence of the correct product attributes, such as a manufacturer's webpage or product photographs.[18]

Amazon's generative-AI listing tool produces titles and product descriptions that sellers can review, edit, and submit as their contributions. These submissions affect what consumers see only when Amazon selects them for display. The tool's effect therefore depends on a seller's ability to influence the displayed content.

In the Entry Data, 90% of entrants with an observed offer history have a mean new-offer count of one over their first 180 days. For these single-listing products, the sole seller is typically the main contributor to the product page and is more likely to hold editorial authority, which creates a direct channel through which the AI tool can affect displayed content. On multiple-listing products, sellers share the product page, and resellers may face constraints on their ability to replace higher-priority contributions. These constraints do not eliminate the AI channel: a brand owner or another contributor with sufficient authority may still use the seller-facing tool to update the shared page.

4.2 Market-Level Treatment Intensity

A natural way to study a generative-AI writing tool is to compare listings with generative-AI markers, or listings with observed adoption of the tool, with those that do not adopt. However, adoption is a seller choice, and sellers who adopt a

[18]See Amazon Seller Forums, "Someone Changed Your Title or Bullets? Understanding Listing Contributions," posted by an Amazon representative, https://sellercentral.amazon.com/seller-forums/discussions/t/fa843ca8-06aa-4138-9fac-8e3ebac06943.

new writing tool may already differ from non-adopters in important ways. They may be better sellers, invest more in their listings, or be experiencing rising or declining sales. This endogenous selection can bias a difference-in-differences estimate even if the parallel pre-trends assumption appears to hold. In addition, competition may affect the adoption decisions and outcomes of other sellers in the same market.

We therefore use market-level ownership structure, measured by the share of single-listing products, as a measure of treatment intensity to the AI tool, where a market is defined by a keyword for incumbents and a leaf category for entrants. For a keyword, SingleShare is the share of products that appeared among its top-three products at any point before the launch that are single-listing. For a category, it is the share of the category's incumbent products that are single-listing. This measure reflects differences in editorial control over listing content, rather than directly measuring brand ownership, tool adoption, or the display of AI-generated text. Treatment intensity thus varies across markets rather than through sellers' adoption choices. Whether a product page carries one offer or multiple offers reflects how the product was brought to market, including brand registration and distribution arrangements.

Specifically, we estimate

$$y_{kt} = \sum_{s \neq t_0} \beta_s \times \left(\text{SingleShare}_k \times \mathbf{1}[t = s]\right) + \mu_k + \tau_t + \varepsilon_{kt}, \tag{2}$$

where y_{kt} is the outcome for market k in period t; SingleShare_k is the market's standardized pre-launch share of single-listing products; t_0 denotes the omitted reference period; and μ_k and τ_t are market and period fixed effects. The coefficients β_s capture differential changes in outcomes associated with a one-standard-deviation increase in the pre-launch single-listing share, relative to t_0. Standard errors are clustered at the market level.

The identifying assumption is that, absent the generative-AI listing tool, markets with higher pre-launch shares of single-listing products would have followed parallel outcome trends. The event studies presented below assess this assumption separately for each outcome. For entry, the pre-launch coefficients are flat. For sales ranks, markets with higher single-listing shares exhibit a pre-trend, so we interpret those estimates cautiously. For click and conversion shares, the sample begins in

April 2023, leaving only five pre-launch months and no pre-period spanning the introduction of ChatGPT.

Difference-in-Differences Specification. We also estimate a two-way fixed-effects difference-in-differences specification using the same continuous treatment-intensity measure:

$$y_{kt} = \beta \times (\text{SingleShare}_k \times \text{Post}_t) + \mu_k + \tau_t + \varepsilon_{kt}, \tag{3}$$

where y_{kt}, SingleShare_k, μ_k, and τ_t are defined as in Equation (2), and Post_t indicates periods after the September 2023 release of the generative-AI listing tool. The coefficient β captures the differential change in the outcome from the pre-launch to the post-launch period associated with a one-standard-deviation increase in the pre-launch single-listing share. Standard errors are clustered at the market level.[19] The results are not sensitive to this difference, consistent with our identification relying on variation in ownership structure that appears to capture differential exposure to Amazon's Generative-AI listing tool rather than to ChatGPT more generally.

We also report estimates for a window that ends before the release of Rufus, Amazon's generative-AI shopping assistant, in February 2024. Later AI tools, including Rufus, Amelia, and ChatGPT Search, may also affect market outcomes, so ending the window in January 2024 allows us to examine the effects of the listing tool before these tools were introduced.

4.3 Ownership Structure Predicts Text Revision and Generative-AI Markers

We test whether text revisions and generative-AI marker use following the release of Amazon's AI listing tool are more prevalent among products with fewer listings using the Wayback Machine Panel. For each product i and quarter t, we construct

[19] In the quarterly regressions, we exclude the fourth quarter of 2022 through the second quarter of 2023, the period between the release of ChatGPT and the launch of Amazon's generative-AI listing tool, so that the pre-period, from the third quarter of 2021 through the third quarter of 2022, is not affected by ChatGPT. This exclusion is not possible for the monthly click and conversion shares, which are available only from April 2023.

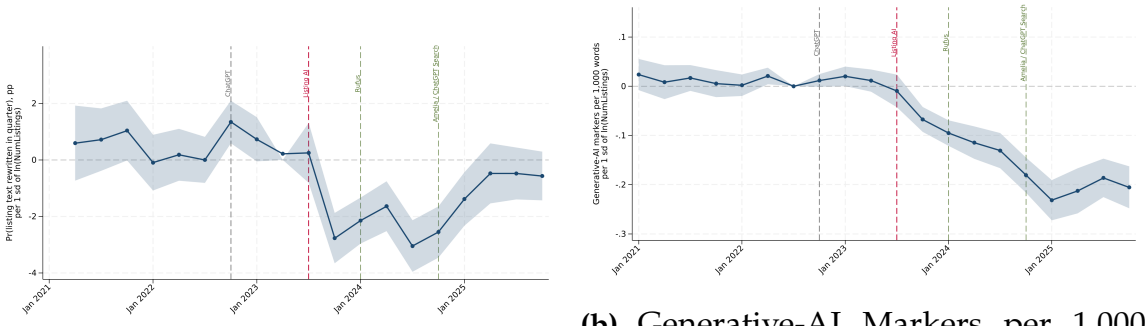
two outcomes: an indicator for whether its title or product description changed between its last capture in that quarter and its most recent preceding capture, and the number of generative-AI marker occurrences per 1,000 words of its title and product description. We define L_i as the standardized log of the product's average number of new-condition listings during the pre-launch period. We estimate:

$$y_{it} = \alpha_i + \gamma_t + \sum_{q \neq q_0} \beta_q \mathbf{1}[t = q] \times L_i + \varepsilon_{it}, \tag{4}$$

where α_i and γ_t are product and quarter fixed effects. Thus, β_q captures how the outcome for products with more listings evolves relative to products with fewer listings, after accounting for product and quarter fixed effects. Standard errors are clustered at the product level, and coefficients are measured relative to q_0, the third quarter of 2022, the last quarter before the introduction of ChatGPT.

Figure 3 plots the estimates from Equation (4). Before the tool's release, the coefficients are generally close to zero. The temporary spike around the introduction of ChatGPT may reflect holiday-season updating by sellers of low-listing products rather than a response to any of the marked events. After the release, the revision coefficient falls to about -3 percentage points per standard deviation of log listings and remains between -2 and -3 percentage points through 2024, indicating that products with fewer listings are revised more often. The coefficient for generative-AI marker use is also close to zero before the release and becomes negative afterward, reaching about -0.1 markers per 1,000 words by early 2024 and about -0.2 from late 2024 onward, relative to a baseline of 0.8 markers per 1,000 words in the third quarter of 2022. This indicates that products with fewer listings increase their use of generative-AI markers more after the release. These patterns support using ownership structure as a measure of differential exposure to Amazon's generative-AI listing tool.

Similar Results Using External ChatGPT Markers. In addition, we use words that prior studies document as overused by ChatGPT (Kobak et al., 2025; Liang et al., 2024), such as *delve*, *showcase*, *intricate*, and *meticulous*. Because these words are selected independently of our data, the results do not depend on how we construct the generative-AI marker set. Within products in the Wayback Machine panel, estimates from Equation (1) show that the share of listings containing these



(a) Probability of Text Revision

(b) Generative-AI Markers per 1,000 Words

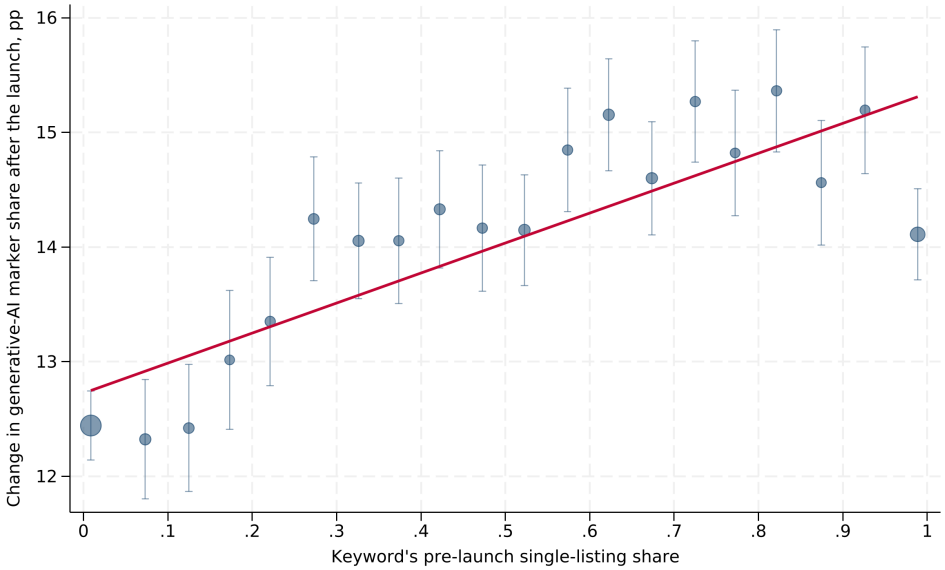
Notes: The figure plots the coefficients β_q from Equation (4), which interact quarter indicators with the standardized log of the product's average daily number of new-condition listings during the pre-launch period. The outcome in Figure 3a is an indicator for whether the product's title or description changed between its last Wayback Machine capture in the quarter and its most recent preceding capture; the first quarter of 2021 is omitted because few products have a preceding capture. The outcome in Figure 3b is the number of generative-AI marker occurrences per 1,000 words of the title and product description, using the 40 words described in Section 3.1. Coefficients are expressed per one-standard-deviation increase in log listings: in percentage points in Figure 3a and in generative-AI marker occurrences per 1,000 words in Figure 3b. Estimates are normalized relative to the third quarter of 2022, the last quarter before the introduction of ChatGPT. Bands show 95% confidence intervals, with standard errors clustered at the product level. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure 3: Text Revision and Generative-AI Markers by Number of Listings

words is 3.1 percentage points higher by the end of 2024 than in the third quarter of 2022, relative to a baseline of 5.5 percent in that quarter. Appendix G shows that, after the release of the generative-AI listing tool, these words are adopted more frequently by products with fewer listings.

Market-Level Marker Adoption. We also check whether the market-level treatment intensity measure tracks the changes in listing text it is intended to capture. For each keyword, we compute the share of clicked products whose listing text contains a Generative-AI marker, separately before and after the launch, and relate the change in this share to the keyword's pre-launch single-listing share. The sample contains 223,261 keywords with at least three clicked products in each period. Figure 4 shows that keywords with no single-listing products experience an increase in marker share of about 12 percentage points. The increase grows with the single-listing share, reaching about 15 percentage points for keywords where 60 percent or more of clicked products are single-listing products. This

pattern supports the use of the pre-launch single-listing share as a measure of differential exposure to the AI listing tool.



Notes: The figure plots the change in the share of a keyword's clicked products whose listing text carries a generative-AI marker, as defined in Section 3.1, against the keyword's pre-launch share of single-listing products. The change is measured as the post-launch marker share minus the pre-launch marker share, in percentage points. The sample includes 223,261 keywords with at least three clicked products in each period. Keywords are grouped into twenty equal-width bins of the single-listing share. Each circle represents a bin mean, with its area proportional to the number of keywords in the bin. Caps show 95% confidence intervals for the bin means, and the line is a bin-count-weighted linear fit.

Figure 4: Generative-AI Marker Adoption and Pre-Launch Single-Listing Share

5 Product Entry and Incumbent Sales

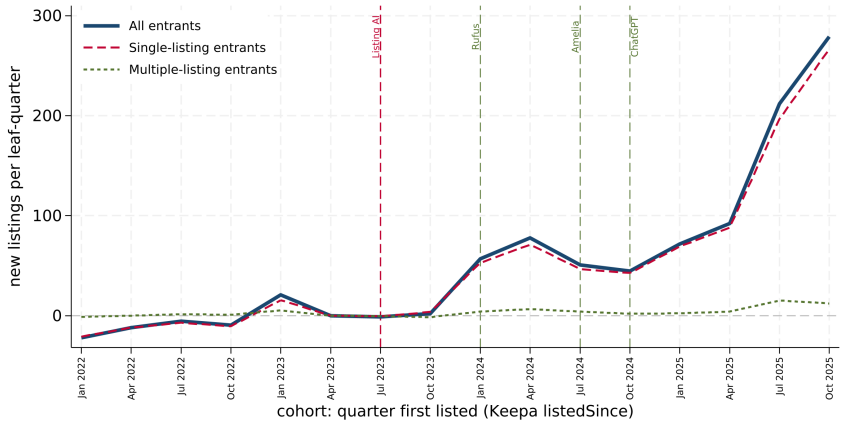
5.1 Entry by Ownership Structure

Generative AI may lower the cost of bringing new products to market by reducing the effort required to create product-listing text. Producing a viable listing no longer requires the same level of copywriting effort, potentially making entry easier for sellers. We estimate an event study at the category–cohort level:

$$N_{\ell c} = \alpha_\ell + \sum_{q \neq q_0} \gamma_q \, \mathbf{1}[c = q] + \varepsilon_{\ell c}, \tag{5}$$

where $N_{\ell c}$ denotes the number of new products listed in product category ℓ and quarterly cohort c, and α_ℓ are category fixed effects that absorb level differences in category size.

Figure 5 plots the estimated cohort coefficients from Equation (5), using q_0, the second quarter of 2023, the last quarter before the release of Amazon's AI listing tool, as the baseline. The estimates capture within-category changes in the number of new listings per quarter relative to that cohort. We find a sharp increase in new product entry, driven almost entirely by single-listing products, for which the editorial-control mechanism described above implies greater treatment intensity to the generative-AI listing tool. The cohort from the second quarter of 2024 contains about 78 more new products per category than the baseline cohort: approximately 71 single-listing products and 7 multiple-listing products. By the second quarter of 2025, the corresponding increases are 92, 88, and 4. The concentration of entry growth among single-listing products is consistent with generative-AI tools lowering the cost of creating product pages.



Notes: The figure plots the estimated cohort coefficients from Equation (5), measured in new products per category. The solid line shows estimates for all entrants, while the dashed and dotted lines show estimates separately for single-listing and multiple-listing entrants. Estimates are normalized relative to the second quarter of 2023, the last quarter before the release of Amazon's generative-AI listing tool. Bands show 95% confidence intervals, with standard errors clustered at the category level. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure 5: Entrants by Ownership Structure

5.2 Entry by Treatment Intensity

We now examine entry using the treatment-intensity design and the Entry Data, defining a market as a product category. Figure 6a plots the coefficients β_c from Equation (2), which interact cohort indicators with the standardized pre-launch share of single-listing products among incumbents in each category. The sample includes 3,642 leaf categories with at least ten incumbents and at least one entrant. The coefficients are flat before the tool's release and rise afterward. Beginning in the first quarter of 2024, a category with a one-standard-deviation higher pre-launch single-listing share experiences roughly 5 to 9 additional new products per quarter relative to the baseline, rising to about 19 to 20 by late 2025.

Panel A of Table 2 reports the difference-in-differences estimate from Equation (3). Under the parallel-trends assumption, a one-standard-deviation increase in the pre-launch single-listing share increases entry by 9.2 products per category per quarter after the launch. Restricting the sample to the period before the release of Rufus yields an increase of 2.8 products, equivalent to 55 percent of the pre-launch mean. Entry thus responds in the months immediately following the launch of the listing tool, before Amazon introduced its demand-side tools. In levels, this early response is about 30 percent of the full-window effect, with most of the increase accumulating over 2024 and 2025.

In the Entry Data, 76 percent of entrants with known seller origin are foreign, defined as products for which most offers come from sellers located in China. Foreign entrants account for most of the increase in entry: before Rufus, 2.3 of the 2.8 additional entrants are foreign, as shown in Appendix E.

5.3 Sales Ranks of Incumbents

More new product introductions are likely to increase market competition, so we examine their effects on incumbent products. In the keyword data, we define incumbents as products that rank among the top three for a given keyword at any point during the sample period. We examine incumbents' sales ranks using Equation (2), where the market k is a keyword and the period t is a quarter, with the third quarter of 2022, the last quarter before the introduction of ChatGPT, as the reference period. The outcome is the average log sales rank, $\ln(\text{SalesRank})$, across incumbent products for each keyword. Because a higher numerical sales

Panel A. Entry: new products per category and quarter

	Number of new products		Relative to pre-launch mean	
	Pre-Rufus	Full	Pre-Rufus	Full
SingleShare × Post	2.7543***	9.1844***	0.5537***	1.1722*
	(0.9826)	(2.4839)	(0.1800)	(0.6168)
Category Fixed Effects	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes
Observations	32,778	61,914	18,666	35,258

Panel B. Keyword level: mean ln *sales rank of the keyword's products*

	All products		Top three		Top one	
	Pre-Rufus	Full	Pre-Rufus	Full	Pre-Rufus	Full
SingleShare × Post	0.0758***	0.0998***	0.0633***	0.0134***	0.0620***	-0.0096***
	(0.0010)	(0.0011)	(0.0014)	(0.0015)	(0.0019)	(0.0020)
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,135,321	6,718,545	2,863,105	6,135,225	2,206,246	4,727,670

Panel C. Product level: ln *sales rank of the individual product*

	Listed before 2022		Listed January 2022 to launch	
	Pre-Rufus	Full	Pre-Rufus	Full
SingleShare × Post	0.0809***	0.1348***	0.1216***	0.1460***
	(0.0010)	(0.0012)	(0.0080)	(0.0089)
Product Fixed Effects	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes
Observations	8,290,681	17,765,744	165,365	429,949

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). Panel A uses the Entry Data, with the outcome defined as the number of new products first listed in a category in a given quarter in columns 1 and 2, and this number divided by the category's pre-launch mean in columns 3 and 4. Panel B uses keyword-level Brand Analytics Data, with the outcome defined as the mean log Keepa sales rank of products appearing in the Brand Analytics Data in columns 1 and 2, the top-three products in columns 3 and 4, and the top product in columns 5 and 6. Panel C uses product-level Brand Analytics Data for products listed before 2022 in columns 1 and 2 and those listed between January 2022 and the launch in columns 3 and 4. The outcome is the log of the product's mean Keepa sales rank in the quarter. Standard errors are clustered at the category, keyword, and product levels in Panels A, B, and C, respectively. The full period runs from the first quarter of 2021 through the fourth quarter of 2025, while the pre-Rufus period ends in the fourth quarter of 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2: Entry and Sales Ranks by Treatment Intensity

rank indicates lower sales, a positive coefficient implies that incumbent sales ranks worsen more in markets with higher treatment intensity.

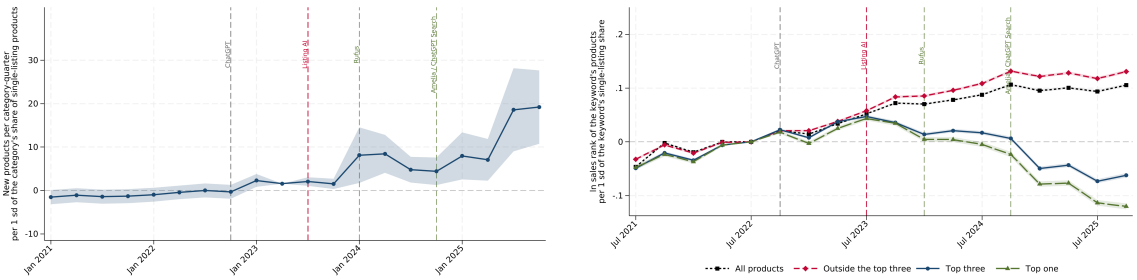
Figure 6b shows the corresponding event study from Equation (2). Immediately after the release of the generative-AI listing tool, the sales ranks of top products and the remaining products begin to diverge. Sales ranks improve for the top

products, with a larger improvement for the top product than for the top three as a group, while those of other products worsen. The gap between the top product and products outside the top three widens from about 1 percentage point at the launch to 25 percentage points by the end of 2025. However, we note the presence of a pre-trend: in keywords with a higher single-listing share, sales ranks were already worsening before the listing tool's release, and this increase continues through the months following the introduction of ChatGPT.

We also examine the difference-in-differences estimates, which should be interpreted with caution because of this pre-trend. Panel B of Table 2 reports the estimates from Equation (3). In columns 1 and 2, a one-standard-deviation increase in the single-listing share is associated with a 7.6 percent increase in sales rank before Rufus and a 10.0 percent increase over the full window. Columns 3 to 6 report the estimates for the keyword's top-three products and top product. For the top three, the estimates are 6.3 percent before Rufus and 1.3 percent over the full window; for the top product, the corresponding estimates are 6.2 percent and -1.0 percent. Given the pre-trend in Figure 6b, these estimates may largely reflect the continuation of trends that began before the listing tool rather than its effect. We therefore emphasize the more informative result: the post-launch divergence between the top products and the rest, shown in Figure 6b.

Panel C of Table 2 reports the product-level evidence, controlling for product fixed effects. We examine products listed before 2022 and products listed between January 2022 and the launch separately, using balanced panels of products with an observed sales rank in every quarter.[20] For products listed before 2022, a one-standard-deviation increase in the single-listing share is associated with an 8.1 percent increase in sales rank before Rufus and a 13.5 percent increase over the full window. For products listed between January 2022 and the launch, the corresponding increases are 12.2 and 14.6 percent.

[20]For the newer products, the pre-period covers only the first three quarters of 2022, as sales ranks are not observed before entry.



(a) Number of Entrants **(b)** Log Sales Rank

Notes: The figure plots the coefficients β_c from Equation (2), using the number of entrants and log sales rank as outcomes in Figure 6a and Figure 6b, respectively. Estimates are normalized relative to the third quarter of 2022, the last quarter before the introduction of ChatGPT. Standard errors are clustered at the category and keyword levels in Figure 6a and Figure 6b, respectively.

Figure 6: Product Entry and Incumbent Sales Ranks

6 Market Concentration

6.1 Click and Conversion Shares of the Top Three Products

Sales rank is a relative measure of sales, so changes in sales rank do not necessarily imply a change in concentration. The Brand Analytics Data report the click and conversion shares of the top-three products for each keyword, allowing us to measure directly how much demand is concentrated among these products. Our concentration measure is therefore defined across products within a keyword, rather than across sellers or firms.

We examine concentration using their combined click and conversion shares, which measure the concentration ratios of the top three products, as outcomes in Equation (2), where the market k is a keyword and the period t is a month. March 2023 and earlier months are excluded because of the Brand Analytics reporting change on March 25, 2023.

Figure 7a and Figure 7b show the estimated changes in click and conversion concentration ratios across all keywords. Both concentration ratios rise significantly after the release of the generative-AI listing tool. Averaged over October 2023 to December 2024, a one-standard-deviation increase in the single-listing share raises the top-three click share by 2.0 percent and the top-three conversion share by 7.6 percent.

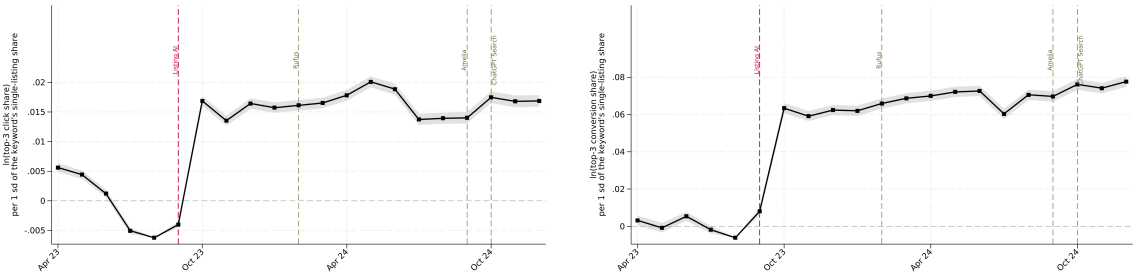
Table 3 reports the difference-in-differences estimates from Equation (3). The

top-three click share increases by 1.1 percent before Rufus and 1.4 percent over the full window, while the corresponding increases in the top-three conversion share are 5.2 and 6.5 percent.

	ln top-3 click share		ln top-3 conversion share	
	Pre-Rufus (1)	Full (2)	Pre-Rufus (3)	Full (4)
SingleShare × Post	0.0108***	0.0140***	0.0524***	0.0653***
	(0.0002)	(0.0002)	(0.0007)	(0.0007)
Keyword Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	4,441,570	9,327,297	3,564,800	7,486,080

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). Columns 1 and 2 use the log combined click share of the top-three products, while columns 3 and 4 use their log combined conversion share, from the Brand Analytics Data. Standard errors are clustered at the keyword level. The sample runs from April 2023 through December 2024, while the pre-Rufus period ends in January 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Concentration by Treatment Intensity



(a) Log Top-Three Click Share **(b)** Log Top-Three Conversion Share

Notes: Figure 7a and Figure 7b plot coefficients from Equation (2) for all keywords. The outcomes are the log monthly top-three click share and conversion share, respectively. The sample covers April 2023 through December 2024; March 2023 and earlier months are excluded because of the Brand Analytics reporting change on March 25, 2023. Estimates are normalized relative to August 2023, the last month before the release of the generative-AI listing tool. Bands show 95% confidence intervals, with standard errors clustered at the keyword level. Vertical lines mark the releases of Amazon's generative-AI listing tool, Rufus, Amelia, and ChatGPT Search.

Figure 7: Click and Conversion Concentration Ratios

6.2 Rewriting Ability and Query Language

Small Effects in Categories with Limited Seller Rewriting. For media products, sellers have limited scope to change product pages because these pages are largely standardized. Books, movies, recorded music, and physical video games are generally listed on product pages constructed from authoritative metadata, such as title, author or artist, International Standard Book Number (ISBN) or European Article Number (EAN), publisher, studio or label, and edition information, rather than from seller-authored text. Third-party sellers typically do not edit these shared product pages. As a result, media products are much less directly exposed to the text-generation channel of Amazon's generative-AI listing tools.[21]

As a robustness check, we estimate Equation (3) separately for media keywords, where sellers have limited ability to rewrite product pages. We find that the estimated effects are much smaller for media keywords, with magnitudes only about one-half to one-sixth of the main estimates in Table 3, suggesting that the main results are unlikely to be driven by broader changes unrelated to listing text. See Appendix F for more details.

Concentration Increases and Returns to Text. We further examine whether the concentration effect varies with query language in a way consistent with a listing-text mechanism. We separate keywords into navigational brand-name searches and other searches. Within non-brand searches, we further distinguish queries by the number of descriptive attributes specified, such as material, color, size, use case, recipient, style, or value. We estimate Equation (3) separately for each group. In Table 4, each row represents a keyword group, and each cell reports the estimated change in the concentration ratio associated with a one-standard-deviation increase in the single-listing share.[22]

Row 1 of Table 4 shows that navigational brand-name searches, for which consumers specify the brand they seek, exhibit only a 0.3 percent increase in

[21] We exclude Kindle and e-book products from the media control group. These products are often created through Kindle Direct Publishing (KDP), where authors have greater control over titles and descriptions and may use generative AI in the publishing process. E-books are therefore at least partly exposed to the treatment and are classified with the treatable categories rather than the media control group.

[22] The Brand Analytics Data identify keywords by their search strings. We examine how the search popularity of different types of keywords changes over this period in Appendix K.

top-three click share before Rufus and a 0.2 percent increase over the full window. In percentage points, the brand-name effect is not statistically significant at the 5 percent level and turns negative over the full window (-0.05, $t = -1.9$). In row 2, all other queries exhibit increases of 1.3 and 1.6 percent, respectively, consistent with concentration increasing through better matching of listing text to the query.

Rows 3 to 6 show that, within non-brand queries, the effect rises with the number of descriptive attributes specified. Before Rufus, queries with no descriptive attributes exhibit a 0.9 percent increase in top-three click share, those with one or two attributes a 1.7 to 1.8 percent increase, and those with three or more a 2.3 percent increase. For conversion shares, the ordering is strictly monotonic. The effect for richly descriptive queries is roughly eight times that for brand-name searches. We interpret these patterns as evidence consistent with a channel operating through listing text.

6.3 Concentration and Turnover by Pre-Launch Dominance

The increase in concentration can reflect changes along two distinct dimensions: the inequality of demand across products and the identity of the products that capture that demand. Concentration may rise because products that were already among the top products before the launch gain further share, or because products that were not previously at the top move into top positions and subsequently capture a disproportionate share of demand. In the latter case, the market can become more concentrated even as the identity of the top products changes.

These possibilities have different implications for how generative AI reshapes competition. If pre-launch top products gain further share, the generative-AI listing tool reinforces existing advantages and moves the market toward a more winner-take-all distribution of demand with little change in the identity of top products. If different products move into top positions and then capture a large share of demand, the market may also become more winner-take-all, but through a reshuffling of which products occupy the top positions. We therefore distinguish changes in the concentration of demand from changes in the persistence and identity of top products.

We examine whether the concentration effect differs between keywords with and without a dominant top product before the launch of the generative-AI listing

	Pre-Rufus		Full window	
	(1) ln click share	(2) ln conv. share	(3) ln click share	(4) ln conv. share
Brand-name query	0.0030***	0.0226***	0.0024***	0.0247***
	(0.0006)	(0.0018)	(0.0007)	(0.0019)
All other (non-brand) queries	0.0129***	0.0584***	0.0162***	0.0728***
	(0.0003)	(0.0008)	(0.0003)	(0.0009)
No attribute	0.0094***	0.0435***	0.0121***	0.0559***
	(0.0003)	(0.0009)	(0.0003)	(0.0010)
1 attribute	0.0175***	0.0755***	0.0223***	0.0919***
	(0.0005)	(0.0016)	(0.0006)	(0.0018)
2 attributes	0.0166***	0.0773***	0.0196***	0.0929***
	(0.0011)	(0.0038)	(0.0012)	(0.0041)
3+ attributes	0.0234***	0.0851***	0.0263***	0.0997***
	(0.0028)	(0.0106)	(0.0031)	(0.0113)
Keyword Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	4,452,670	3,570,160	9,350,607	7,497,336

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). Columns 1 and 3 use the log combined click share of the top-three products as the outcome, while columns 2 and 4 use their log combined conversion share, from the Brand Analytics Data. Rows 1 and 2 report estimates for navigational brand-name queries and all other queries, respectively. Rows 3 to 6 further separate non-brand queries by the number of descriptive attributes specified in the search string: zero, one, two, and three or more. These attributes include material, color, size, use case, recipient, style, or value. Standard errors are clustered at the keyword level. The sample runs from April 2023 through December 2024, while the pre-Rufus period ends in January 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Heterogeneous Concentration Effects by Query Language

tool. We measure dominance using the margin of the top product over the second-ranked product, averaged over the weeks before the launch in the Brand Analytics Data. We split keywords at the median of this margin and estimate Equation (3) separately for keywords with and without a dominant top product. In Appendix H, we also use alternative measures of dominance based on the average click share of the top product, its tenure at the top, the number of distinct products that are most clicked in at least one week, and top-three turnover, and find similar results.

We next decompose the top-three shares into the portion held by the keyword's pre-launch top products, defined as products that appear among the top-three products in at least half of the weeks before the launch, and the portion held by other products. Table 5 reports estimates of Equation (3) for the log total top-three share and for each component.

Columns 1 and 4 of Table 5 show that the concentration effect arises primarily in keywords without a dominant top product. Before Rufus, a one-standard-deviation increase in the single-listing share raises the top-three click share by 1.5 percent in keywords without a dominant top product, compared with a decline of 0.4 percent in keywords with one. The corresponding estimates for the top-three conversion share are 6.1 and 1.8 percent.

Columns 2, 3, 5, and 6 show where this change comes from. Across all keywords, the part of the top-three click share held by pre-launch top products falls by 2.3 percent before Rufus, while the part held by other products rises by 2.6 percent. The same pattern appears in keywords with a dominant top product, where the part held by pre-launch top products falls by 2.4 percent and the part held by other products rises by 3.1 percent. In keywords without a dominant top product, the part held by pre-launch top products changes little, falling by 0.9 percent for the click share and rising by 0.2 percent for the conversion share, while the part held by other products rises by 3.0 and 6.3 percent, respectively.

Panel B of Table 5 examines turnover among the top-three products. We measure top-three turnover as the share of a week's top-three products that were not among the top three in the previous week, and top-product turnover as the share of weeks in which the top product differs from that in the previous week, both averaged over the month and expressed in percentage points. In keywords with a dominant top product, the top product changes more often after the launch: a one-standard-deviation increase in the single-listing share raises top-product turnover by 1.6 percentage points before Rufus. In keywords without a dominant top product, top-three turnover falls by 0.3 percentage points, indicating that the set of top products becomes slightly more stable. These changes are small relative to the average weekly top-three turnover of 31 percent and average top-product turnover of 35 percent.

Taken together, these results suggest that the generative-AI listing tool increases concentration not by reinforcing the position of dominant pre-launch products. In fact, pre-launch top products lose share, while other products gain share and move into top positions. Where a product already dominated a keyword before the launch, concentration changes little. Instead, the increase in concentration is concentrated in markets that were more competitive before the launch, where demand shifts toward a smaller set of products that were not previously among

the top products. The generative-AI listing tool therefore appears to weaken the persistence of pre-existing advantages while making demand more concentrated among the products that reach the top positions.

	ln top-3 click share			ln top-3 conversion share		
	Total	Pre-launch top products	Other products	Total	Pre-launch top products	Other products
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Log top-three shares						
Pre-Rufus						
All keywords	0.0107***	-0.0232***	0.0260***	0.0523***	-0.0155***	0.0551***
	(0.0002)	(0.0007)	(0.0008)	(0.0007)	(0.0010)	(0.0014)
Keywords with a dominant top product	-0.0041***	-0.0242***	0.0314***	0.0183***	-0.0204***	0.0462***
	(0.0003)	(0.0008)	(0.0013)	(0.0008)	(0.0011)	(0.0021)
Keywords without a dominant top product	0.0154***	-0.0088***	0.0295***	0.0609***	0.0023	0.0632***
	(0.0004)	(0.0016)	(0.0010)	(0.0012)	(0.0021)	(0.0019)
Full window						
All keywords	0.0138***	-0.0289***	0.0201***	0.0652***	-0.0203***	0.0585***
	(0.0002)	(0.0008)	(0.0008)	(0.0007)	(0.0011)	(0.0014)
Keywords with a dominant top product	-0.0045***	-0.0300***	0.0299***	0.0237***	-0.0274***	0.0511***
	(0.0003)	(0.0009)	(0.0014)	(0.0009)	(0.0012)	(0.0021)
Keywords without a dominant top product	0.0179***	-0.0162***	0.0280***	0.0715***	-0.0007	0.0704***
	(0.0004)	(0.0017)	(0.0010)	(0.0012)	(0.0022)	(0.0019)
Panel B. Turnover among the top-three products (percentage points)						
	Top-three turnover	Change of top product				
Pre-Rufus						
All keywords	-0.171***	0.481***				
	(0.016)	(0.033)				
Keywords with a dominant top product	-0.086***	1.607***				
	(0.021)	(0.043)				
Keywords without a dominant top product	-0.275***	0.700***				
	(0.026)	(0.051)				
Full window						
All keywords	-0.319***	0.465***				
	(0.016)	(0.031)				
Keywords with a dominant top product	-0.076***	1.991***				
	(0.020)	(0.040)				
Keywords without a dominant top product	-0.443***	0.734***				
	(0.026)	(0.047)				
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). In Panel A, columns 1 and 4 use the log combined click share and conversion share of the top-three products. Columns 2 and 5 use the log of the portion of these shares held by pre-launch top products, defined as products appearing among the top-three products in at least half of the weeks before the launch, while columns 3 and 6 use the log of the portion held by other products. Keyword-months in which a component is zero are excluded. In Panel B, the outcomes are the share of a week's top-three products that were not among the top three in the previous week and the share of weeks in which the top product differs from that in the previous week, averaged over the month and expressed in percentage points. A keyword has a dominant top product if the pre-launch margin of the top product over the second-ranked product is above the median. Standard errors are clustered at the keyword level. The full period runs from April 2023 through December 2024, while the pre-Rufus period ends in January 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Concentration and Turnover by Pre-Launch Dominance

7 Listing Text and the Basis of Competition

Concentration increases in markets with greater treatment intensity even as these markets experience more entry. We next examine why by studying how listing text changes and how demand reallocates across products. We then relate these shifts to the prices and ratings of the products receiving the most demand. As competing listings become more similar, text may become less useful for distinguishing among products, shifting competition toward other, non-text attributes.

7.1 How Listing Text Changes

We examine how listing text changes in keywords with greater treatment intensity. We use products in the Wayback Machine Panel that are captured both before the introduction of ChatGPT and after the release of the tool, with the product description captured in both periods. For each product, we compute the change between its last capture through the third quarter of 2022 and its last capture from the fourth quarter of 2023 onward, and average these changes across products within each keyword with at least three such products. We then estimate

$$\Delta y_k = \beta \, \text{SingleShare}_k + \lambda_{t_k} + \varepsilon_k, \tag{6}$$

where Δy_k is the average change in outcome y across the products of keyword k, SingleShare_k is the keyword's standardized single-listing share, and λ_{t_k} are quarter fixed effects for the timing of the keyword's captures, defined by the median quarter of its products' captures before the introduction of ChatGPT and after the release of the tool.

Columns 1 and 2 of Table 6 use the change in the number of generative-AI markers per 1,000 words and in the log number of words. Column 3 uses the change in the share of titles that repeat a word, excluding common words such as "for" and "and", three or more times, which is prohibited under Amazon's title rules. Column 4 measures how closely listings resemble the keyword's pre-launch top product, defined as the product that is most clicked in the largest number of weeks before the launch. We use the top product's pre-ChatGPT text as a fixed reference and compute, for each other product, the change in cosine similarity between its title and product description and this reference text. Column 5 uses

the change in peer-text similarity among the keyword's products, multiplied by 100. Peer-text similarity is the mean pairwise cosine similarity across products within a keyword, based on ℓ_2-normalized hashed term-frequency vectors of their titles and product descriptions.

We find that listings adopt more generative-AI text in keywords with greater treatment intensity. A one-standard-deviation increase in the single-listing share is associated with an additional increase of 0.26 generative-AI markers per 1,000 words. Listings in keywords with a higher single-listing share also lengthen less, by 1.2 log points per standard deviation.

Listings in more treated keywords also become less repetitive and more similar to the pre-launch top product. Amazon's title rules prohibit repeating a word more than twice, a practice sellers may use to match additional search terms. Before the introduction of ChatGPT, 17 percent of titles repeat a word three or more times, and a one-standard-deviation increase in the single-listing share reduces this share by 2.9 percentage points, about one-sixth of its pre-ChatGPT level. A one-standard-deviation increase in the single-listing share also raises similarity to the top product's pre-ChatGPT text by 0.47 points.[23] These patterns suggest that the generative-AI listing tool improves some observable dimensions of listing quality and makes sellers' listings more similar to the pre-launch top product.

Peer-text similarity also rises slightly more in keywords with a higher single-listing share, by 0.06 points per standard deviation. This change is small relative to the average level of similarity, which is about 29 points on the same scale.

The second and third panels of Table 6 split products by seller origin. They average product-level changes within each keyword over products sold by domestic and foreign sellers, respectively, restricting the sample to keywords with at least three such products. We define foreign sellers as sellers located in China and domestic sellers as sellers outside China, most of whom are located in the United States. The adoption of generative-AI text and convergence toward the pre-launch top product are concentrated among domestic sellers. In keywords with a higher single-listing share, listings of domestic sellers add 0.65 generative-AI markers per 1,000 words per standard deviation and move 0.98 points closer to the top product's pre-ChatGPT text. In contrast, listings of foreign sellers show slightly

[23] Using titles only, the estimate is 0.47 points. Using product descriptions only, the estimate is 0.48 points.

fewer markers and move 0.23 points away from the top product's text. Titles that repeat a word three or more times decline in both groups. The largest changes in observed listing text occur among domestic sellers, whereas entry and demand gains are concentrated among foreign sellers. This contrast suggests that the benefits of generative AI extend beyond the sellers whose own listings change most, consistent with demand shifting toward other product attributes, including price.

	Markers per 1,000 words (1)	100 × ln words (2)	Titles repeating a word 3+ times (%) (3)	Similarity to top product ×100 (4)	Peer-text similarity ×100 (5)
All sellers					
SingleShare	0.258***	-1.176***	-2.948***	0.474***	0.056***
	(0.012)	(0.080)	(0.056)	(0.030)	(0.020)
Observations	167,851	167,851	167,851	94,355	167,851
Domestic sellers					
SingleShare	0.654***	1.220***	-2.029***	0.977***	0.263***
	(0.019)	(0.125)	(0.081)	(0.043)	(0.029)
Observations	111,144	111,144	111,144	63,557	111,144
Foreign sellers					
SingleShare	-0.146***	-0.256**	-1.712***	-0.230***	-0.005
	(0.026)	(0.116)	(0.184)	(0.041)	(0.027)
Observations	37,690	37,690	37,690	19,835	37,690
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes

Notes: The table reports estimates of Equation (6) using products in the Wayback Machine Panel observed with product descriptions both before ChatGPT and after the launch. Columns 1 to 5 use changes in generative-AI markers per 1,000 words, log listing length, title repetition, similarity to the pre-launch top product, and peer-text similarity, respectively; the last three outcomes are multiplied by 100 where applicable. The first panel uses all products, while the second and third panels average changes over domestic and foreign sellers, respectively. Foreign sellers are products for which at least half of active new-condition offers between March 25 and September 9, 2023 come from sellers located in China. Robust standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: Changes in Listing Text by Treatment Intensity

7.2 Non-Text Attributes Gain Importance

To understand which products gain after the release of the generative-AI listing tool, we use all keyword-product pairs in the Brand Analytics Data and estimate

$$y_{ikt} = \sum_j \left[\delta_j X_{ijt} \times \text{SingleShare}_k \times \text{Post}_t + \gamma_j X_{ijt} \times \text{Post}_t + \theta_j X_{ijt}\right] + \beta\, \text{SingleShare}_k \times \text{Post}_t + \alpha_{ik} + \tau_t + \varepsilon_{ikt}. \tag{7}$$

where y_{ikt} is either an indicator, multiplied by 100, for product i being among the top-three products for keyword k, or the log of the product's click or conversion share; X_{ijt} is product characteristic j in period t; α_{ik} are keyword-product pair fixed effects; and τ_t are time fixed effects. The coefficient δ_j measures how the relationship between characteristic j and product performance changes after the launch with the keyword's treatment intensity. Identification comes from changes in the characteristics and performance of the same product within a keyword over time. We estimate the specification using products in the Brand Analytics Data for product characteristics and the Wayback Machine sample for changes in listing text. Standard errors are clustered at the keyword level.

Table 7 reports the estimates. Columns 1 to 3 use products in the Brand Analytics Data. In keywords with higher treatment intensity, products sold by foreign sellers gain 1.5 and 2.4 percent in click and conversion share per standard deviation of the single-listing share. Lower prices predict demand gains only among foreign products: a one-standard-deviation lower log price raises their click and conversion shares by 0.5 and 1.0 percent, whereas higher-priced domestic products gain 0.5 and 0.7 percent. Seller reputation matters for both groups: products sold by sellers with ratings of at least 95 percent gain 0.57 percentage points in top-three probability and 2.1 and 3.1 percent in click and conversion share. Product ratings, however, have opposite associations across seller origins. Highly rated domestic products gain 0.13 percentage points in top-three probability and 0.6 percent in conversion share, whereas highly rated foreign products lose 0.50 percentage points and 1.0 and 1.3 percent in top-three probability, click share, and conversion share, respectively. Products with few reviews and older products, in contrast, lose.

Returns to Text Decline as Text Becomes More Similar. Columns 4 to 6 use the Wayback Machine sample. Products whose listings become more similar to other listings in the same keyword lose: a one-standard-deviation larger increase in peer-text similarity lowers top-three probability by 0.11 percentage points and click and conversion shares by 0.7 and 1.2 percent. Products whose titles become more repetitive also lose, with a 0.29 percentage point decline in top-three probability. On the other hand, products whose listings add more generative-AI language or become longer gain slightly.

These patterns suggest that as listing text becomes more similar across competing products, the return to differentiating through text declines. They are consistent with listing text playing a larger role for domestic sellers, while lower prices become more important for foreign products.

7.3 Ratings and Prices of Top Products

Ultimately, it is important to understand how the generative-AI listing tool affects consumers. Most product characteristics in Table 7 are predetermined or difficult to adjust in the short run, whereas ratings and prices can respond more directly to market conditions. Ratings may reflect both seller investments in product quality and how well the information in the listing matches consumers' expectations. Prices, in turn, are chosen by sellers and respond to competitive conditions. In Appendix K, we show that consumer search behavior also changes over this period. Searches become less subjective in their wording: consumers search less using style and value terms such as "cute" or "best", while searches using brand names increase through 2024. Searches containing product specifications remain relatively stable. Consumers also use fewer queries specifying multiple attributes, such as material, size, or use case, particularly in 2025.

We therefore examine how the characteristics of products clicked by consumers change after the launch. For each keyword and quarter, we compute the mean Keepa star rating and mean log price for products in the Brand Analytics Data, the top-three products, and the top product. We then estimate Equation (3) with keyword and quarter fixed effects for all products and separately for domestic and foreign products.

Table 8 reports the estimates. Ratings rise after the launch. Over the full window,

	Brand Analytics products			Wayback Machine Panel		
	Top three (pp) (1)	ln click share (2)	ln conversion share (3)	Top three (pp) (4)	ln click share (5)	ln conversion share (6)
Foreign × SingleShare × Post	0.28***	0.0153***	0.0235***			
	(0.09)	(0.0019)	(0.0023)			
Domestic: Log price × SingleShare × Post	0.62***	0.0052***	0.0071***			
	(0.04)	(0.0009)	(0.0013)			
Foreign: Log price × SingleShare × Post	-0.15**	-0.0050***	-0.0102***			
	(0.07)	(0.0014)	(0.0018)			
Domestic: High rating × SingleShare × Post	0.13**	0.0011	0.0062***			
	(0.07)	(0.0013)	(0.0016)			
Foreign: High rating × SingleShare × Post	-0.50***	-0.0100***	-0.0131***			
	(0.08)	(0.0015)	(0.0018)			
Low reviews × SingleShare × Post	-0.37***	-0.0055***	-0.0049***			
	(0.06)	(0.0011)	(0.0013)			
Seller rating $\geq$ 95% × SingleShare × Post	0.57***	0.0210***	0.0305***			
	(0.07)	(0.0016)	(0.0021)			
Product age × SingleShare × Post	-1.13***	-0.0098***	-0.0071***			
	(0.04)	(0.0008)	(0.0010)			
Δ Peer-text similarity × SingleShare × Post				-0.11***	-0.0070***	-0.0115***
				(0.04)	(0.0011)	(0.0013)
Δ Generative-AI markers × SingleShare × Post				0.08	0.0034***	0.0031**
				(0.05)	(0.0012)	(0.0014)
Δ ln words × SingleShare × Post				0.28***	0.0065***	0.0091***
				(0.06)	(0.0014)	(0.0015)
Δ Title repetition × SingleShare × Post				-0.29***	-0.0024*	-0.0029**
				(0.06)	(0.0012)	(0.0014)
Pair fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27,416,604	33,729,220	24,337,516	6,750,992	8,311,121	7,098,027

Notes: The table reports estimates of Equation (7) with keyword-product pair and time fixed effects. Each row reports the listed characteristic interacted with SingleShare × Post; the characteristic itself, its interaction with Post, and SingleShare × Post are included but not reported. Columns 1 and 4 use top-three status, multiplied by 100, while columns 2, 3, 5, and 6 use log click and conversion shares. Columns 1 to 3 use current product characteristics entered jointly; seller rating and product age are common across seller types, while log price and high rating, defined as a rating of at least 4.5 stars, enter separately for domestic and foreign sellers. Columns 4 to 6 use standardized changes in listing text from the Wayback Machine Panel. Foreign sellers are sellers located in China. Standard errors are clustered at the keyword level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Changing Returns to Product Characteristics and Listing Text

a one-standard-deviation increase in the single-listing share is associated with a 0.008-star increase in the average rating of the top-three products. This increase is concentrated among domestic products. Their ratings rise across specifications, including by 0.006 stars among the top-three products over the full window, while ratings of foreign products remain approximately unchanged.

Prices also fall after the launch in markets with higher treatment intensity. Over the full window, prices fall by 1.7 percent for products in the Brand Analytics Data, 2.3 percent for the top-three products, and 2.8 percent for the top product. Because the price of the top product is already declining slightly before the launch, we

interpret this estimate with caution. Two patterns contribute to the decline. First, foreign sellers gain top positions: their share among the top-three products rises by 3.6 percentage points, and Table 7 shows that lower-priced foreign products are more likely to gain. This increase comes increasingly from new products over time. Over the full window, foreign products first listed after the launch gain 4.3 percentage points of the top three, while older foreign products lose 0.7 percentage points, as shown in Appendix I. The increase in the foreign share is also larger in keywords where listing text converges more after the launch; see Appendix J. Second, average prices among domestic top-three products fall by 0.8 percent, while prices of foreign top-three products remain approximately unchanged.

Overall, the products receiving the most demand have lower average prices and slightly higher average ratings in markets with higher treatment intensity. This pattern is consistent with the greater importance of non-text attributes documented in Table 7 as listing text becomes less useful for distinguishing among products. Greater concentration therefore coincides with a shift toward products with more favorable observable attributes, although these outcomes do not by themselves identify changes in consumer welfare.

8 Conclusion

This paper studies how generative AI affects product entry and market concentration through its use in product-listing text. We examine Amazon's September 2023 rollout of seller-facing generative-AI listing tools, combining data on product entry, consumer clicks and purchases, product characteristics, and historical listing text. Our identification exploits an institutional feature of Amazon that creates differential ability to revise listing text using the generative-AI listing tool. We measure this treatment intensity using markets' pre-launch shares of single-listing products, whose sellers have greater control over the shared product page. This market-level design allows us to study differential exposure to the tool without relying directly on sellers' endogenous adoption decisions.

We find that markets with greater treatment intensity experience both more product entry and greater concentration of consumer demand among the top three products. The concentration effects are substantially smaller in media categories, where sellers have limited ability to rewrite listing text, and are larger for more

	Brand Analytics products		Top-three products		Top product	
	Pre-Rufus	Full	Pre-Rufus	Full	Pre-Rufus	Full
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Star rating						
All products	0.0021***	0.0081***	0.0018***	0.0080***	0.0012***	0.0071***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Domestic products	0.0018***	0.0060***	0.0016***	0.0056***	0.0015***	0.0055***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Foreign products	-0.0027***	-0.0010**	-0.0014**	0.0011*	-0.0011	0.0011
	(0.0005)	(0.0004)	(0.0007)	(0.0006)	(0.0008)	(0.0007)
Observations: all	2,885,967	6,184,215	2,526,811	5,414,595	2,253,825	4,829,625
Observations: domestic	1,973,531	4,228,995	1,383,200	2,964,000	1,148,868	2,461,860
Observations: foreign	1,033,179	2,213,955	580,517	1,243,965	470,344	1,007,880
Panel B. ln *price*						
All products	-0.0081***	-0.0172***	-0.0117***	-0.0233***	-0.0123***	-0.0276***
	(0.0004)	(0.0004)	(0.0005)	(0.0005)	(0.0006)	(0.0006)
Domestic products	-0.0057***	-0.0131***	0.0026***	-0.0076***	0.0013	-0.0119***
	(0.0006)	(0.0006)	(0.0008)	(0.0008)	(0.0010)	(0.0010)
Foreign products	0.0007	-0.0006	-0.0018	-0.0028**	-0.0009	-0.0013
	(0.0008)	(0.0008)	(0.0012)	(0.0012)	(0.0014)	(0.0015)
Observations: all	2,857,960	6,124,200	2,510,536	5,379,720	2,234,057	4,787,265
Observations: domestic	1,972,915	4,227,675	1,382,430	2,962,350	1,145,529	2,454,705
Observations: foreign	1,035,370	2,218,650	581,462	1,245,990	471,177	1,009,665
Panel C. Foreign share of products (percent)						
SingleShare × Post	3.39***	3.29***	3.32***	3.60***	3.34***	3.89***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
Observations	2,728,334	5,846,430	2,327,759	4,988,055	2,017,813	4,323,885
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports estimates from Equation (3). Panels A and B use, respectively, the mean star rating and mean log price of products in the Brand Analytics Data, measured using Keepa data. Within Panels A and B, rows report results for all products, domestic products, and foreign products. Panel C uses the share of products sold by foreign sellers, multiplied by 100. Foreign products are those with at least half of their offers in the quarter from sellers located in China. Columns 1 and 2 use products in the Brand Analytics Data, columns 3 and 4 use the keyword's top-three products, and columns 5 and 6 use its top product. Standard errors are clustered at the keyword level. The full sample period runs from the third quarter of 2021 through the fourth quarter of 2025, excluding the fourth quarter of 2022 through the second quarter of 2023, while the pre-Rufus period ends in the fourth quarter of 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Post-Launch Changes in Ratings and Prices

descriptive search queries. These patterns suggest that the increase in concentration is closely related to the listing-text channel. We also find suggestive evidence that, as generative-AI text becomes more prevalent and competing listings become more

similar, non-text attributes become increasingly important for holding a top-three position, particularly ratings, review counts, and prices.

Finally, the increase in concentration is not accompanied by worse observable outcomes for consumers. In more exposed markets, the products receiving the most clicks have lower average prices and slightly higher average ratings. These patterns are consistent with demand shifting toward products with stronger non-text attributes as listing text becomes less useful for distinguishing among alternatives. Taken together, our results show that competition through AI-generated text can reshape market structure in ways that go beyond lowering entry costs: it can increase product entry while simultaneously concentrating demand and shifting competition toward non-text product attributes.

References

Luis Aguiar and Joel Waldfogel. Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music. *Journal of Political Economy*, 126(2):492–524, 2018.

Simon P. Anderson and Régis Renault. Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model. *RAND Journal of Economics*, 30 (4):719–735, 1999.

Tania Babina, Anastassia Fedyk, Alex He, and James Hodson. Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics*, 151:103745, 2024.

Sagit Bar-Gill, Erik Brynjolfsson, and Nir Hak. Helping Small Businesses Become More Data-Driven: A Field Experiment on eBay. *Management Science*, 70(11): 7345–7372, 2024. doi: 10.1287/mnsc.2021.02026.

Heski Bar-Isaac, Guillermo Caruana, and Vicente Cuñat. Search, Design, and Market Structure. *American Economic Review*, 102(2):1140–1160, 2012.

Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at Work. *Quarterly Journal of Economics*, 140(2):889–942, 2025.

Junhui Jeff Cai, Xian Gu, Liugang Sheng, Mengjia Xia, Linda Zhao, and Wu Zhu. AI as "Co-founder": GenAI for Entrepreneurship, 2025. arXiv:2512.06506.

Jingyi Cui, Gabriel Dias, and Justin Ye. Signaling in the Age of AI: Evidence from Cover Letters. Working paper, arXiv:2509.25054, 2025. URL https://arxiv.org/abs/2509.25054.

Kevin Zheyuan Cui, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Management Science*, 2026. doi: 10.1287/mnsc.2025.00535. Articles in Advance. Published online February 27, 2026.

Dirk Czarnitzki, Gastón P. Fernández, and Christian Rammer. Artificial Intelligence and Firm-Level Productivity. *Journal of Economic Behavior & Organization*, 211: 188–205, 2023.

Anil R. Doshi and Oliver P. Hauser. Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content. *Science Advances*, 10(28): eadn5290, 2024.

Michael Ewens, Ramana Nanda, and Matthew Rhodes-Kropf. Cost of Experimentation and the Evolution of Venture Capital. *Journal of Financial Economics*, 128(3): 422–442, 2018.

Lu Fang, Zhe Yuan, Kaifu Zhang, Dante Donati, and Miklos Sarvary. Generative AI and Firm Productivity: Field Experiments in Online Retail. Working paper, arXiv:2510.12049v1, October 2025, 2025. URL https://arxiv.org/abs/2510.12049v1.

Anais Galdin and Jesse Silbert. Making Talk Cheap: Generative AI and Labor Market Signaling. Working paper, arXiv:2511.08785, 2025. URL https://arxiv.org/abs/2511.08785.

Samuel Goldberg and H. Tai Lam. Generative AI & Creative Goods: Market Expansion, Crowd-Out, and Copyright. *Quarterly Journal of Economics*, 2026. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5152649. Forthcoming.

Xiang Hui, Oren Reshef, and Luofeng Zhou. The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. *Organization Science*, 35(6):1977–1989, 2024. doi: 10.1287/orsc.2023.18441.

Sueyoul Kim, Ginger Zhe Jin, and Eungik Lee. Does Generative AI Crowd Out Human Creators? Evidence from Pixiv. NBER Working Paper 34733, National Bureau of Economic Research, 2026. URL https://www.nber.org/papers/w34733.

Dmitry Kobak, Rita González-Márquez, Emőke-Ágnes Horvát, and Jan Lause. Delving into LLM-Assisted Writing in Biomedical Publications through Excess Vocabulary. *Science Advances*, 11(27):eadt3813, 2025. doi: 10.1126/sciadv.adt3813.

Weixin Liang, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, Haotian Ye, Sheng Liu, Zhi Huang, Daniel A. McFarland, and James Y. Zou. Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 29575–29620. PMLR, 2024. URL https://proceedings.mlr.press/v235/liang24b.html.

Yao Lu, Gordon M. Phillips, and Jia Yang. The Impact of Cloud Computing and AI on Industry Dynamics and Concentration. NBER Working Paper 32811, National Bureau of Economic Research, 2024.

Kristina McElheran, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. AI Adoption in America: Who, What, and Where. *Journal of Economics & Management Strategy*, 33(2):375–415, 2024.

Shakked Noy and Whitney Zhang. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654):187–192, 2023.

Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. arXiv:2302.06590, 2023.

Michael Spence. Job Market Signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973.

Ole Teutloff, Johanna Einsiedler, Otto Kässi, Fabian Braesemann, Pamela Mishkin, and R. Maria del Rio-Chanona. Winners and Losers of Generative AI: Early Evidence of Shifts in Freelancer Demand. *Journal of Economic Behavior & Organization*, 235:106845, 2025. doi: 10.1016/j.jebo.2024.106845.

Joel Waldfogel. How Digitization Has Created a Golden Age of Music, Movies, Books, and Television. *Journal of Economic Perspectives*, 31(3):195–214, 2017.

Emma Wiles, Zanele Munyikwa, and John J. Horton. Algorithmic Writing Assistance on Jobseekers' Resumes Increases Hires. *Management Science*, 71(12): 10144–10164, 2025. doi: 10.1287/mnsc.2024.04528.

Online Appendix

The appendix is intended for online publication.

Contents of Appendix

A Amazon's Generative-AI Tools for Sellers **45**

B Data Collection Procedure **48**

C The Wayback Machine Panel **49**

D Generative-AI Marker Frequency and Selection **49**

D.1 Robustness of Generative-AI Marker Selection 50

E Entry by Seller Origin **51**

F Media Categories with Limited Seller Rewriting **52**

F.1 Smaller Concentration Effects in Media Categories 53

G Robustness Using Words Overused by ChatGPT **55**

H Alternative Measures of Pre-Launch Dominance **55**

I New and Older Products Among Top Products **56**

J Foreign Share and Text Convergence **57**

K Changes in Consumer Search Queries **58**

A Amazon's Generative-AI Tools for Sellers

We describe several seller-facing generative-AI tools that complement Amazon's listing text tool, introduced on September 13, 2023.[24] These tools support the

[24] See https://www.aboutamazon.com/news/small-business/amazon-sellers-generative-ai-tool.

creation of multiple listings, access to business information, the production of advertising and branded content, and updates to existing listings. Table A1 summarizes the main developments from September 2023 through May 2025.

Amazon first expanded the types of information sellers could use to generate listing content. In March 2024, it introduced the option to provide a URL from the seller's own website, allowing the tool to extract product information and generate content for an Amazon listing. By that time, sellers could also generate listing content from a product image, in addition to a short text description.[25] In September 2024, Amazon announced a bulk workflow that would generate titles, bullet points, and descriptions for multiple products from a spreadsheet containing limited product information. By May 2025, Amazon reported that this workflow had been rolled out.[26]

In September 2024, Amazon also introduced Project Amelia, a generative-AI assistant integrated into Seller Central. Initially released in beta to a subset of U.S. sellers, the assistant answers questions about selling on Amazon and retrieves information specific to a seller's business. For example, sellers can request summaries of sales, units sold, and customer traffic, compare performance with the previous year, and ask follow-up questions about particular products. At launch, Amazon also described plans to extend the assistant to diagnosing operational problems and, in some cases, taking actions on sellers' behalf.[27]

During the same month, Amazon introduced generative-AI support for A+ Content and announced Video Generator. A+ Content allows brands to add enhanced text and visual material to product detail pages. Its generative-AI features help sellers produce marketing text and product imagery, which they can refine through additional prompts. Video Generator creates video advertisements from a single product image; its initial beta was planned for selected U.S. advertisers in the fourth quarter of 2024. These tools are intended to reduce the time and resources required to produce marketing content, including for sellers without dedicated creative teams.[28]

[25] See https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-powered-product-listings.

[26] See https://www.aboutamazon.com/news/innovation-at-amazon/amazon-generative-ai-seller-growth-shopping-experience, published September 19, 2024 and updated May 8, 2025.

[27] See https://www.aboutamazon.com/news/innovation-at-amazon/amazon-project-amelia.

[28] See https://sellercentral.amazon.com/seller-forums/discussions/t/6179a8be-23e0-48be-9200-cc583e465de5 and https://advertising.amazon.com/library/news/ai-video-

In May 2025, Amazon introduced Enhance My Listing, a tool for maintaining existing listings. Built using Amazon Bedrock, it draws on customer activity and engagement data to recommend changes to titles, descriptions, and product attributes, including missing details. Sellers can review, customize, accept, or decline the recommendations. The tool initially began rolling out to selected U.S. sellers, expanding Amazon's support for ongoing listing updates.[29]

Date	Tool or development	Capability
Sep 13, 2023	Listing text generation	Generates titles, descriptions, and listing details from a short text input.
Mar 2024	URL-based listing creation	Generates listing content from a product page on the seller's own website.
Sep 2024	Bulk listing workflow announced	Generates content for multiple products from a spreadsheet containing limited information.
Sep 2024	Project Amelia beta	Answers seller questions and retrieves business metrics through Seller Central.
Sep 2024	Generative AI for A+ Content	Helps brands create marketing text and imagery for product detail pages.
Sep 2024	Video Generator announced	Creates video advertisements from a product image; initial beta planned for Q4 2024.
May 2025	Enhance My Listing	Recommends revisions and missing information for existing listings, subject to seller review.

Notes: Compiled from the Amazon announcements and seller-guidance pages cited in the text. Dates refer to announcements or initial releases and do not necessarily indicate availability to all sellers.

Table A1: Rollout of Amazon's Seller-Facing Generative-AI Tools

`generator-live-image`.

[29] See the May 2025 update to the announcement cited in footnote 26.

B Data Collection Procedure

To start, we identify three primary product categories with a relatively high number of new products: Kitchen & Dining, Home Decor, and Office Products. Within each primary category, we delve into precisely defined product categories (third-level subcategories) such as "Mechanical Pencils", "Stemmed Water Glasses", and "Kids' Lunch Boxes." We randomly select 250 distinct subcategories for data collection.

Round 1 For each of the 250 categories, we collect data on the top 1,000 products with the highest 30-day average sales rank. Additionally, we include a random sample comprising 3% of products from the top 10,000 most popular products within each category. We also collect data on all new products introduced within the last two months.[30]

Round 2 and Onwards Each round of data collection spans approximately three days. In each subsequent round t, we track existing products on the list, $\mathbb{S}_{t-1}$, from the previous round, and dynamically include new products:

1. Check keepa.com to identify newly introduced products in the 250 categories, denoted as $\mathbb{E}_t$.
2. Update the list of products that are continuously monitored to include these new products: $\mathbb{S}_t = \mathbb{S}_{t-1} \cup \mathbb{E}_t$.
3. Collect information on $\mathbb{S}_t$.
4. Move onto the next period and start from step 1, with $t = t + 1$.

As a result, the number of products we monitor gradually grows over rounds, enabling us to effectively track information regarding product entry and exit, as well as pricing and sales data. We implemented this procedure from May 2023 through November 2024, gathering information on approximately 3.1 million products from more than 160,000 independent third-party sellers, of which about 47 percent were first listed after the September 2023 launch of Amazon's generative-AI listing tools.

[30] In rare cases where the number of new products in the last two months exceeds 1,000, we focus on the top 1,000 products in this subset.

C The Wayback Machine Panel

Table A2 summarizes the Wayback Machine Panel used for the text analyses, including capture frequency, coverage around the launch, and text length. Non-media products have a median of 10 captures, with a median interval of three weeks between captures, and 69% have captures both before and after the September 2023 launch. Media products have fewer captures, with a median of 7, and much shorter text: the median title contains 37 characters, and product descriptions are typically absent. Within the sample, multiple-listing products are captured more frequently than single-listing products, with medians of 13 and 7 captures, respectively. Single-listing products, however, have longer product descriptions: the median product description contains 835 characters, compared with 688 characters for multiple-listing products.

	All	Non-media	Media	By ownership structure	
				Single-listing	Multiple-listing
Products	94,317	81,745	12,572	32,210	56,294
Captures per product (median)	9	10	7	7	13
Captures per product (mean)	37.5	40.5	18.1	21.3	48.7
Days between captures (median)	21.5	21	24	29	18
Share with captures before and after launch	0.669	0.688	0.546	0.558	0.761
First capture (median year)	2022	2022	2022	2022	2022
Last capture (median year)	2025	2025	2025	2025	2025
Title length, characters (median)	122	133	37	137	117
Description length, characters (median)	722	813	0	835	688

Notes: The table reports summary statistics for the Wayback Machine Panel, which contains historical captures of Amazon product pages for products in the Brand Analytics Data from January 2020 through August 2026. Columns 2 and 3 separate non-media and media products. Columns 4 and 5 separate single-listing and multiple-listing products, where the number of listings is measured as the average daily number of new-condition listings during the pre-launch period; products without listing information are excluded from these columns. The launch refers to the September 13, 2023 release of Amazon's generative-AI listing tools. Text length is measured using the product title and description.

Table A2: Summary Statistics of the Wayback Machine Panel

D Generative-AI Marker Frequency and Selection

Table A3 reports the 49 words selected from listing revisions, as described in Section 3.1, and shows how frequently the 40 retained generative-AI markers appear in the Brand Analytics Data and the Entry Data. A listing is classified as

containing generative-AI markers if its title or product description contains at least one of these 40 words.

Word	Ratio	Brand Analytics Data (%)	Entry Data (%)	Word	Ratio	Brand Analytics Data (%)	Entry Data (%)
captivating	17.11	0.53	0.52	integrates	4.52	0.21	0.37
pfas†	11.54	–	–	shared†	4.38	–	–
sophistication	8.37	0.99	2.19	gatherings	4.28	0.96	4.46
catering	6.86	0.35	1.36	elegance	4.26	2.39	6.19
inviting	6.67	0.40	0.89	uninterrupted	4.25	0.26	0.26
ensuring	6.61	7.12	15.42	elevate	4.18	2.13	3.76
seeking	6.43	0.76	0.96	offering	4.15	3.47	5.42
setups†	6.38	–	–	households†	3.96	–	–
cater	6.32	0.27	0.37	ample	3.92	1.66	2.82
meticulously	6.28	0.82	1.10	effortless	3.86	2.61	8.03
enduring	5.83	0.36	0.46	differ†	3.81	–	–
boasting	5.76	0.24	0.30	practicality	3.75	0.68	3.16
cherished	5.67	0.46	0.97	enthusiasts	3.72	1.37	3.60
und†	5.56	–	–	thoughtful	3.69	2.14	6.79
hosting	5.40	0.26	1.02	preserving	3.67	0.28	1.30
simplifies	5.35	0.18	1.03	crafted	3.63	8.39	16.35
inspires	5.21	0.13	0.10	targeted	3.60	0.31	0.04
savor	5.03	0.19	0.26	effortlessly	3.58	2.73	5.96
embrace	4.93	0.89	0.62	feline†	3.56	–	–
sessions	4.91	0.79	1.05	maximizing	3.39	0.23	0.71
multitasking	4.89	0.13	0.07	robust	3.33	1.30	2.54
delightful	4.85	0.85	1.50	unforgettable	3.33	0.67	0.67
enhancing	4.82	1.45	4.41	apartments†	3.31	–	–
tackles	4.75	0.07	0.04	mornings†	3.22	–	–
appeal	4.52	1.00	2.57				
Any word		26.15	45.80				

Notes: The table reports the 49 words whose share among words added in listing revisions is at least three times higher after the release of ChatGPT than before in both random halves of non-media products in the Wayback Machine Panel, and that are added at least 100 times after the release. The ratio equals each word's share among added words after the release of ChatGPT divided by its share before the release. Words marked with † are excluded from the generative-AI marker set. For the remaining 40 words, the last two columns report the share of products whose title or product description contains each word, separately for the Brand Analytics Data and the Entry Data. The last row reports the share of products whose text contains at least one of these 40 words.

Table A3: Selection and Frequency of Generative-AI Markers

D.1 Robustness of Generative-AI Marker Selection

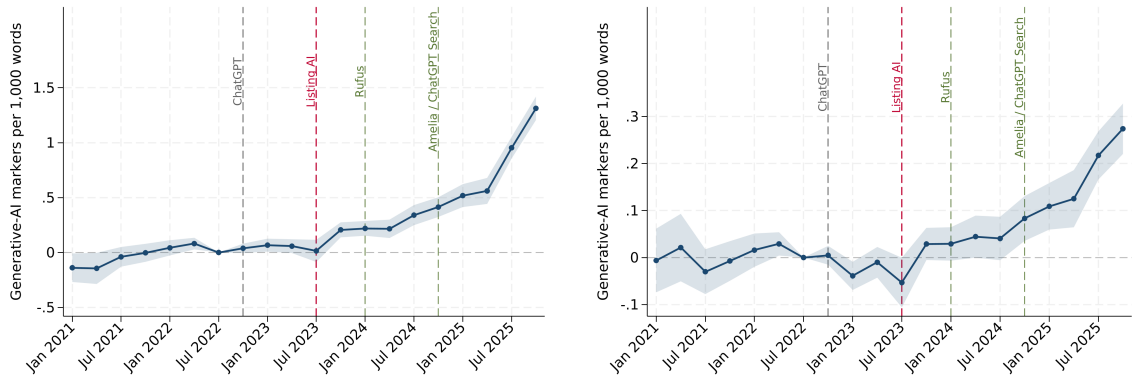
We examine whether the within-product increase in generative-AI markers depends on how the marker words are selected. We use two alternative word lists selected based on dash use rather than timing. The first consists of the 70 words that

are at least twice as common in Entry Data listings with dashes as in listings without dashes, comparing listings of similar length and without removing any words. The second is constructed from the Wayback Machine Panel. We randomly split products into two halves. Within each half, we apply the same dash-based procedure to non-media products, comparing listings with and without dashes of similar length, and retain every word that appears in at least 200 listings and is at least twice as common in listings with dashes. We then measure each product's text using the words selected from the other half, so that no product's text is used to select the words applied to that product. For each product and quarter, we use the number of occurrences of these words per 1,000 words of the title and product description as the outcome in Equation (1).

Figure A1a shows that, using the 70 dash-based words, the coefficients are close to zero before the introduction of ChatGPT and remain close to zero until the release of the generative-AI listing tool. The coefficient then rises to 0.2 occurrences per 1,000 words by the end of 2023, 0.4 by the end of 2024, and 1.3 by the end of 2025, relative to a baseline of 10.7 occurrences per 1,000 words in the third quarter of 2022. Figure A1b shows a similar pattern using the words selected from the Wayback Machine Panel. The coefficients are close to zero before the release of the generative-AI listing tool and then rise to 0.08 occurrences per 1,000 words by the end of 2024 and 0.27 by the end of 2025, relative to a baseline of 2.5 occurrences per 1,000 words in the third quarter of 2022.

E Entry by Seller Origin

Table A4 reports the estimates from Panel A of Table 2 separately for domestic and foreign entrants. We classify an entrant as foreign if at least half of its observed offers come from sellers located in China, and as domestic otherwise. Seller origin is available for about 90 percent of entrants listed through the third quarter of 2024 but for few entrants thereafter, so we end the longer window in the third quarter of 2024. Foreign entrants account for most of the increase in entry. Before Rufus, a one-standard-deviation increase in the pre-launch single-listing share is associated with 2.3 additional foreign entrants and 0.2 additional domestic entrants per quarter; through the third quarter of 2024, the corresponding estimates are 4.8 and 0.5. Relative to their pre-launch means, domestic entry rises by 26 to 30



(a) Dash-Based Words from the Entry Data

(b) Words Selected from the Wayback Machine Panel

Notes: The figure plots the coefficients γ_q from Equation (1), using the number of generative-AI marker occurrences per 1,000 words of the title and product description as the outcome. Figure A1a uses the 70 words that are at least twice as common in Entry Data listings with dashes as in listings without dashes. Figure A1b uses words selected from a random half of Wayback Machine products and applied to the other half. Estimates are normalized relative to the third quarter of 2022, the last quarter before the introduction of ChatGPT. Bands show 95% confidence intervals, with standard errors clustered at the product level. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure A1: Within-Product Changes in Generative-AI Markers

percent.

F Media Categories with Limited Seller Rewriting

Amazon's generative-AI listing tools operate through a clear channel: they lower the cost of creating and improving product-listing text. Their effects should therefore be stronger for products whose listing text can be freely edited by sellers and weaker for products whose content is more tightly controlled. At the product level, single-listing products provide a measure of greater exposure to this channel. At the market level, media products provide a useful comparison because sellers have much less ability to rewrite their product pages.

Books, movies, recorded music, and physical video games are generally listed on standardized product pages constructed from authoritative metadata, such as title, author or artist, International Standard Book Number (ISBN) or European Article Number (EAN), publisher, studio or label, and edition information, rather

	Number of new products		Relative to pre-launch mean	
	Pre-Rufus	Through 2024Q3	Pre-Rufus	Through 2024Q3
	(1)	(2)	(3)	(4)
All entrants	2.7543***	5.9548***	0.5537***	0.9111**
	(0.9826)	(1.8490)	(0.1800)	(0.4205)
Domestic entrants	0.2460	0.5224**	0.2572***	0.2998***
	(0.2329)	(0.2612)	(0.0878)	(0.0845)
Foreign entrants	2.2876***	4.7508***	0.5849*	0.8566
	(0.8509)	(1.5435)	(0.3349)	(1.2455)
Category Fixed Effects	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes
Observations: all	32,778	43,704	18,666	24,888
Observations: domestic	32,778	43,704	13,824	18,432
Observations: foreign	32,778	43,704	14,940	19,920

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3), using the Entry Data. Columns 1 and 2 use the number of new products first listed in a category in a given quarter, and columns 3 and 4 divide this number by the category's pre-launch mean. Foreign entrants are products with at least half of their observed offers from sellers located in China. Standard errors are clustered at the category level. Columns 1 and 3 end in the fourth quarter of 2023, while columns 2 and 4 end in the third quarter of 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: Entry by Treatment Intensity and Seller Origin

than seller-authored text. Third-party sellers typically do not edit these shared product pages. As a result, media products are much less directly exposed to the text-generation channel of Amazon's generative-AI listing tools. Even sellers that adopt Amazon's AI listing generator therefore have limited scope to alter media-product pages because much of the catalog information is centrally controlled. We use media keywords as a market-level robustness check for whether our estimated concentration effects operate through listing text. The prevalence of single-listing products is also very similar across media and non-media keywords: the average single-listing share is 42 percent for media keywords and 44 percent for non-media keywords.

F.1 Smaller Concentration Effects in Media Categories

Table A5 reports estimates of Equation (3) for the top-three click and conversion shares, separately for non-media and media keywords. The effects are much

smaller for media keywords. Before Rufus, a one-standard-deviation increase in the single-listing share raises the top-three click share by 1.1 percent in non-media keywords and 0.2 percent in media keywords, and the top-three conversion share by 5.1 and 1.6 percent, respectively. Over the full window, the corresponding estimates are 1.4 and 0.6 percent for the click share and 6.3 and 1.9 percent for the conversion share. These patterns suggest that the concentration effects in Table 3 operate primarily through listing text rather than through changes common to all keywords.

	Non-media keywords		Media keywords	
	Pre-Rufus	Full	Pre-Rufus	Full
	(1)	(2)	(3)	(4)
Panel A. ln *top-3 click share*				
SingleShare × Post	0.0112***	0.0141***	0.0019***	0.0064***
	(0.0002)	(0.0003)	(0.0007)	(0.0007)
Keyword Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	4,160,090	8,736,189	281,580	591,318
Panel B. ln *top-3 conversion share*				
SingleShare × Post	0.0509***	0.0632***	0.0162***	0.0189***
	(0.0007)	(0.0008)	(0.0023)	(0.0024)
Keyword Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
Observations	3,342,120	7,018,452	222,680	467,628

Notes: The table reports the coefficients on SingleShare × Post from regressions based on Equation (3), estimated separately for non-media and media keywords. Panels A and B use the log combined click share and conversion share of the top-three products, respectively, from the Brand Analytics Data. Standard errors are clustered at the keyword level. The sample runs from April 2023 through December 2024, while the pre-Rufus period ends in January 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Concentration by Treatment Intensity: Media and Non-Media Keywords

G Robustness Using Words Overused by ChatGPT

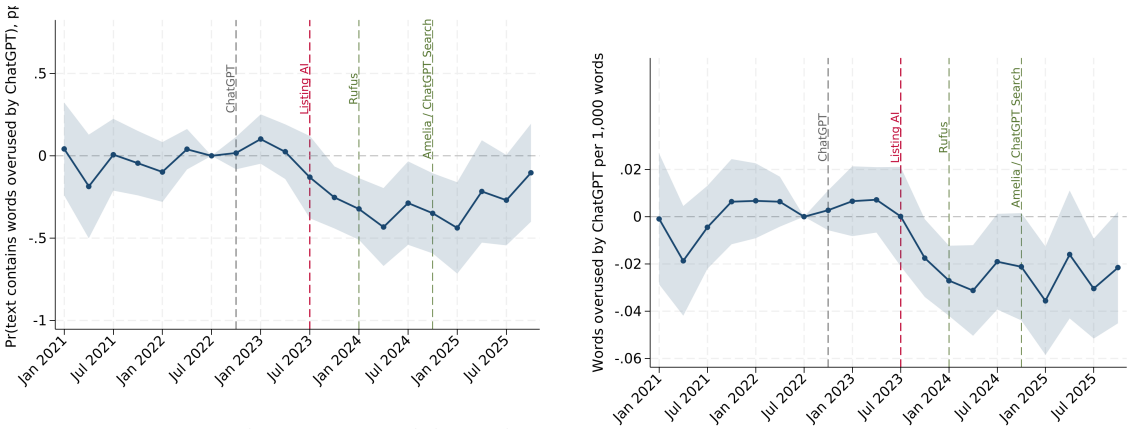
The generative-AI markers used in the paper are selected from our Amazon text data. As a robustness check, we use words that prior studies document as overused by ChatGPT (Kobak et al., 2025; Liang et al., 2024): *delve*, *showcase*, *underscore*, *intricate*, *pivotal*, *meticulous*, *commendable*, *realm*, *boast*, *seamless*, *elevate*, *vibrant*, *tapestry*, *embark*, *testament*, *unparalleled*, *unleash*, *bustling*, and *captivating*. Because these words are selected independently of our data, the results do not depend on how we construct the generative-AI marker set.

We estimate Equation (4) using two outcomes: an indicator for whether the product's title or product description contains at least one of these words, and the number of occurrences of these words per 1,000 words. As in Equation (1), we allow the product fixed effects to differ by whether the full product description is captured to account for systematic differences associated with missing product descriptions.

Figure A2a shows that, before the release of the generative-AI listing tool, the coefficients are close to zero and not statistically significant, including after the introduction of ChatGPT. After the release, the coefficient falls to -0.25 percentage points per standard deviation of log listings by the end of 2023 and remains between -0.3 and -0.4 percentage points through 2024, relative to a baseline of 5.5 percent of listings containing these words in the third quarter of 2022. The coefficient returns toward zero by the end of 2025. Figure A2b shows a similar pattern for the number of occurrences. These patterns indicate that products with fewer listings increase their use of words overused by ChatGPT more after the release of the generative-AI listing tool.

H Alternative Measures of Pre-Launch Dominance

Table A6 examines the analysis in Section 6.3 using five measures of whether a keyword had a dominant top product before the launch: the average click share of the top product, its margin over the second-ranked product, its tenure at the top, the number of distinct products that are top products in at least one week, and top-three turnover. The results are consistent across measures. The concentration effect arises almost entirely in keywords without a dominant top product. Before



(a) Contains Words Overused by ChatGPT

(b) Occurrences per 1,000 Words

Notes: The figure plots the coefficients β_q from Equation (4), which interact quarter indicators with the standardized log of the product's average daily number of new-condition listings during the pre-launch period. The outcome in Figure A2a is an indicator for whether the product's title or product description contains at least one word overused by ChatGPT, and the outcome in Figure A2b is the number of occurrences of these words per 1,000 words. Coefficients are expressed per one-standard-deviation increase in log listings and are normalized relative to the third quarter of 2022, the last quarter before the introduction of ChatGPT. Bands show 95% confidence intervals, with standard errors clustered at the product level. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure A2: Words Overused by ChatGPT and Number of Listings

Rufus, the estimated increase in the top-three click share ranges from -0.6 to 0.2 percent in keywords with a dominant top product, compared with about 1.5 percent in keywords without one.

I New and Older Products Among Top Products

Table A7 decomposes the change in the foreign share in Panel C of Table 8 by product age. For each keyword and quarter, we compute the shares of products that are foreign and new, foreign and older, domestic and new, and domestic and older, where new products are those first listed on Amazon after the launch of the generative-AI listing tool. These four shares sum to 100, so the coefficients within each column sum to zero. Before Rufus, the increase in the foreign share comes mainly from older foreign products, whose share of the top-three products rises by 2.6 percentage points, while new foreign products gain 0.7 percentage points. Over the full window, new foreign products gain 4.3 percentage points of the top

	ln top-3 click share				ln top-3 conversion share			
	Keywords with a dominant top product		Keywords without a dominant top product		Keywords with a dominant top product		Keywords without a dominant top product	
	Pre-Rufus (1)	Full (2)	Pre-Rufus (3)	Full (4)	Pre-Rufus (5)	Full (6)	Pre-Rufus (7)	Full (8)
Top product's click share	-0.0057***	-0.0070***	0.0154***	0.0179***	0.0120***	0.0149***	0.0613***	0.0712***
	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0008)	(0.0009)	(0.0012)	(0.0013)
Margin over second product	-0.0041***	-0.0045***	0.0154***	0.0179***	0.0183***	0.0237***	0.0609***	0.0715***
	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0008)	(0.0009)	(0.0012)	(0.0012)
Tenure of top product	0.0021***	0.0032***	0.0146***	0.0184***	0.0378***	0.0478***	0.0561***	0.0686***
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0009)	(0.0010)	(0.0010)	(0.0011)
Number of distinct top products	0.0000	0.0004	0.0153***	0.0189***	0.0310***	0.0399***	0.0588***	0.0691***
	(0.0002)	(0.0003)	(0.0004)	(0.0004)	(0.0008)	(0.0009)	(0.0013)	(0.0013)
Top-three turnover	0.0014***	0.0010***	0.0155***	0.0213***	0.0288***	0.0359***	0.0657***	0.0789***
	(0.0003)	(0.0003)	(0.0004)	(0.0004)	(0.0008)	(0.0009)	(0.0012)	(0.0013)
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). Columns 1 to 4 use the log combined click share of the top-three products, and columns 5 to 8 use their log combined conversion share, from the Brand Analytics Data. Each row splits keywords at the median of a measure of pre-launch dominance. A keyword has a dominant top product if the top product's average click share, its margin over the second-ranked product, or its tenure is above the median, or if the number of distinct products that are top products in at least one week or top-three turnover is below the median. Standard errors are clustered at the keyword level. The full period runs from April 2023 through December 2024, while the pre-Rufus period ends in January 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Alternative Measures of Pre-Launch Dominance

three, while older foreign products lose 0.7 percentage points. Domestic products lose share in both windows, among both new and older products.

Panels B and C repeat the analysis separately for keywords without and with a dominant top product before the launch, defined as in Section 6.3. Over the full window, new foreign products gain more of the top three in keywords without a dominant top product, by 4.5 percentage points compared with 3.3 percentage points in keywords with one. Older foreign products lose 1.6 percentage points in keywords without a dominant top product but retain their share in keywords with one. In contrast, the decline among older domestic products is larger in keywords with a dominant top product, at 2.0 percentage points compared with 1.0 percentage point in keywords without one.

J Foreign Share and Text Convergence

Table A8 examines whether the increase in the foreign share documented in Panel C of Table 8 is larger in keywords where listing text converges more after the launch. We measure convergence as the change in peer-text similarity among the keyword's products in the Wayback Machine Panel from before ChatGPT to after the launch, as in column 5 of Table 6. Panel A splits keywords at the median change, while Panel B interacts SingleShare × Post with the standardized

change in similarity. The foreign share rises more in keywords where listing text converges more. Over the full window, a one-standard-deviation increase in the single-listing share raises the foreign share of the top-three products by 4.9 percentage points in keywords with an above-median increase in similarity, compared with 3.8 percentage points in keywords with a below-median increase. In Panel B, a one-standard-deviation larger increase in similarity raises the effect by about 0.4 percentage points across outcomes.

K Changes in Consumer Search Queries

The Brand Analytics Data report a search frequency rank for each keyword and week, which ranks keywords by how often consumers search for them. We measure the search popularity of keyword k in quarter t, y_{kt}, as the negative of its mean log search frequency rank in that quarter. We use the 450,588 keywords observed in every quarter from the fourth quarter of 2021 through the fourth quarter of 2025.

We classify keywords according to the words they contain. Table A9 summarizes the definition of each query type, with example keywords and the share of keywords in each type. A keyword can belong to more than one type.

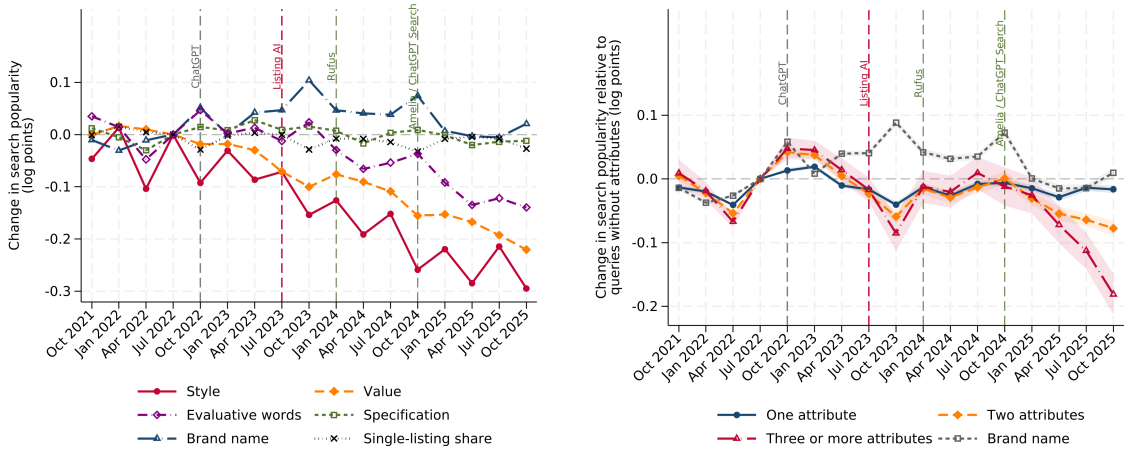
We estimate

$$y_{kt} = \sum_{g \in G} \sum_{q \neq q_0} \theta_{gq} \left(D_{gk} \times \mathbf{1}[t = q] \right) + \mu_k + \tau_t + \varepsilon_{kt}, \tag{A1}$$

where D_{gk} indicates that keyword k belongs to group g, μ_k are keyword fixed effects, and τ_t are quarter fixed effects. The reference period q_0 is the third quarter of 2022, the last quarter before the introduction of ChatGPT. The coefficient θ_{gq} measures the change in the search popularity of group g in quarter q relative to the omitted group. For Figure A3a, we estimate Equation (A1) separately for each query type, so that G contains a single type and the omitted group consists of all other keywords. We also estimate the same specification with D_{gk} replaced by the keyword's standardized single-listing share. For Figure A3b, G contains indicators for keywords with one, two, and three or more attributes, as well as brand-name queries, with queries without attributes as the omitted group. Standard errors are clustered at the keyword level.

Figure A3a shows that consumers increasingly move away from searches containing style and value words. Before the introduction of ChatGPT, the coefficients for value queries are close to zero, while those for style queries fluctuate between -0.10 and 0.01. Style queries then become less popular relative to other queries, with the decline growing after the release of the generative-AI listing tool and reaching about 0.30 log points by the end of 2025. Value queries and queries with evaluative words follow similar patterns, declining by about 0.22 and 0.14 log points, respectively. In contrast, the popularity of specification queries changes little. Brand-name queries become more popular after the introduction of ChatGPT, with the largest increase, about 0.10 log points, occurring in the quarter of the generative-AI listing tool's release, before returning to their earlier level in 2025. These changes begin before the release of the generative-AI listing tool and are not systematically related to our measure of treatment intensity. The search popularity of keywords with a higher single-listing share remains within about 0.03 log points of its pre-ChatGPT level throughout the period.

Figure A3b groups keywords by the number of attributes and plots the corresponding coefficients from Equation (A1), with queries without attributes as the omitted group and brand-name queries included separately. Queries with more attributes become less popular over time. Queries with one attribute change little. By the end of 2025, queries with two attributes decline by about 0.08 log points, while queries with three or more attributes decline by about 0.18 log points, with most of the decline occurring in 2025.



(a) By Query Type **(b)** By Number of Attributes

Notes: Figure A3a plots the coefficients θ_{gq} from Equation (A1), estimated separately for each keyword group and for the standardized single-listing share. Figure A3b plots the coefficients from a single regression based on Equation (A1), by number of attributes and relative to queries without attributes. The outcome is the negative of the keyword's mean log search frequency rank in the quarter. Coefficients are normalized relative to the third quarter of 2022, the last quarter before the introduction of ChatGPT. Standard errors are clustered at the keyword level. The confidence intervals in Figure A3a are omitted because the standard errors are small. In Figure A3b, bands show 95% confidence intervals. Vertical lines mark the releases of ChatGPT, Amazon's generative-AI listing tool, Rufus, and Amelia / ChatGPT Search.

Figure A3: Search Popularity by Query Type and Number of Attributes

	Brand Analytics products		Top-three products		Top product	
	Pre-Rufus	Full	Pre-Rufus	Full	Pre-Rufus	Full
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Shares by seller origin and product age						
Foreign, new	1.25***	4.79***	0.73***	4.27***	0.76***	4.24***
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.03)
Foreign, older	2.10***	-1.52***	2.57***	-0.69***	2.55***	-0.37***
	(0.02)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
Domestic, new	-0.37***	-1.94***	-0.24***	-1.68***	-0.19***	-1.45***
	(0.01)	(0.02)	(0.01)	(0.02)	(0.01)	(0.02)
Domestic, older	-2.99***	-1.33***	-3.06***	-1.89***	-3.12***	-2.42***
	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)	(0.05)
Observations	2,724,078	5,837,310	2,320,276	4,972,020	2,010,743	4,308,735
Panel B. Keywords without a dominant top product						
Foreign, new	1.24***	4.74***	0.77***	4.47***	0.84***	4.54***
	(0.01)	(0.03)	(0.02)	(0.03)	(0.02)	(0.04)
Foreign, older	1.63***	-2.32***	2.29***	-1.60***	2.20***	-1.40***
	(0.04)	(0.04)	(0.05)	(0.06)	(0.06)	(0.07)
Domestic, new	-0.46***	-2.22***	-0.27***	-1.83***	-0.30***	-1.81***
	(0.01)	(0.03)	(0.02)	(0.03)	(0.02)	(0.04)
Domestic, older	-2.41***	-0.21***	-2.79***	-1.04***	-2.75***	-1.33***
	(0.04)	(0.05)	(0.06)	(0.06)	(0.06)	(0.07)
Observations	1,365,896	2,926,920	1,117,361	2,394,345	1,004,437	2,152,365
Panel C. Keywords with a dominant top product						
Foreign, new	1.07***	3.97***	0.60***	3.33***	0.47***	2.90***
	(0.01)	(0.03)	(0.01)	(0.03)	(0.02)	(0.03)
Foreign, older	2.13***	-0.63***	2.26***	0.07	2.10***	0.53***
	(0.04)	(0.04)	(0.05)	(0.05)	(0.06)	(0.06)
Domestic, new	-0.25***	-1.53***	-0.17***	-1.40***	-0.13***	-1.19***
	(0.01)	(0.02)	(0.01)	(0.03)	(0.02)	(0.03)
Domestic, older	-2.96***	-1.81***	-2.69***	-1.99***	-2.44***	-2.24***
	(0.04)	(0.04)	(0.05)	(0.05)	(0.06)	(0.07)
Observations	1,355,634	2,904,930	1,201,263	2,574,135	1,005,249	2,154,105
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the coefficients on SingleShare × Post from separate regressions based on Equation (3). The outcome is the share, multiplied by 100, of products in each group among products in the Brand Analytics Data in columns 1 and 2, the top-three products in columns 3 and 4, and the top product in columns 5 and 6. New products are those first listed in or after the third quarter of 2023. Panels B and C split keywords at the median pre-launch margin of the top product over the second-ranked product. Foreign products are those with at least half of their offers in the quarter from sellers located in China. Standard errors are clustered at the keyword level. The full sample period runs from the third quarter of 2021 through the fourth quarter of 2025, excluding the fourth quarter of 2022 through the second quarter of 2023, while the pre-Rufus period ends in the fourth quarter of 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Foreign and Domestic Shares by Product Age and Pre-Launch Dominance

	Brand Analytics products		Top-three products		Top product	
	Pre-Rufus	Full	Pre-Rufus	Full	Pre-Rufus	Full
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. SingleShare × Post by change in peer-text similarity						
Below-median change	3.87***	3.72***	3.68***	3.82***	3.77***	4.11***
	(0.06)	(0.06)	(0.08)	(0.08)	(0.10)	(0.10)
Above-median change	4.95***	4.77***	4.60***	4.91***	4.76***	5.43***
	(0.06)	(0.07)	(0.09)	(0.09)	(0.11)	(0.11)
Observations: below median	559,321	1,198,545	497,539	1,066,155	433,076	928,020
Observations: above median	563,528	1,207,560	502,040	1,075,800	436,793	935,985
Panel B. Interaction with the change in peer-text similarity						
SingleShare × Post	4.40***	4.24***	4.13***	4.36***	4.26***	4.77***
	(0.04)	(0.05)	(0.06)	(0.06)	(0.07)	(0.08)
SingleShare × Post × Δ Similarity	0.45***	0.39***	0.43***	0.44***	0.42***	0.45***
	(0.04)	(0.05)	(0.06)	(0.06)	(0.07)	(0.07)
Observations	1,122,849	2,406,105	999,579	2,141,955	869,869	1,864,005
Keyword Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports estimates based on Equation (3). The outcome is the share, multiplied by 100, of foreign products among products in the Brand Analytics Data in columns 1 and 2, the top-three products in columns 3 and 4, and the top product in columns 5 and 6. Foreign products are those with at least half of their offers in the quarter from sellers located in China. The change in peer-text similarity is measured at the keyword level using the Wayback Machine Panel. Panel A splits keywords at the median change, while Panel B interacts the standardized change with SingleShare × Post and also includes its interaction with Post. The sample includes keywords for which the change in peer-text similarity can be measured. Standard errors are clustered at the keyword level. The full sample period runs from the third quarter of 2021 through the fourth quarter of 2025, excluding the fourth quarter of 2022 through the second quarter of 2023, while the pre-Rufus period ends in the fourth quarter of 2023. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: Foreign Share by Change in Peer-Text Similarity

Query type	Definition	Example keywords	Share (%)
Style	Contains a word describing appearance or taste, such as "cute", "boho", "vintage", "aesthetic", "rustic", "farmhouse", "luxury", "minimalist", or "decor"	cute beanie; boho canvas wall art; farmhouse valance	1.4
Value	Contains a word describing quality or price, such as "best", "cheap", "affordable", "premium", "professional", "heavy duty", "bulk", or "high quality"	best kitchen knife; cheap bed frame queen; napkins bulk	1.2
Evaluative words	Contains an evaluative word, such as "best", "top rated", "cheap", "good", "cute", "pretty", "cool", "unique", "quality", "perfect", or "amazing"	best kitchen knife; cute weekender bag; cool action figures	0.9
Specification	Contains a material ("stainless", "cotton", "wooden"), color ("black", "clear"), size ("large", "king", "mini"), technical feature ("wireless", "usb", "waterproof"), unit or count ("oz", "inch", "pack"), or a number	wooden foot stool; 8mm wrench; white led rock lights	22.4
Brand name	Contains the brand of one of the keyword's top-three products before the launch	xero shoes; coffee mate powdered creamer; victorinox spartan	17.5
Number of attributes	Number of matches, in keywords that are not brand-name queries, to the style, value, and specification words above or to words describing a use case ("for kitchen", "replacement") or recipient ("for women", "gift"); a number counts once	small freezer chest (one); black floating shelves for bathroom (two); pink king size comforter set (three)	22.6 / 5.2 / 0.7

Notes: The table reports the words used to classify keywords into query types, example keywords, and the share of the keywords in each type. For the number of attributes, the reported shares correspond to keywords with one, two, and three or more attributes.

Table A9: Definitions of Query Types