Pricing Patterns on Dual-Mode E-Commerce Platforms*

Luise Eisfeld†
University of Lausanne and Swiss Finance Institute

Jun Yan‡
Xiamen University

Li Yu§
Liaoning University

This paper is preliminary and is updated frequently.
Comments are welcome.

June 11, 2026

Abstract

Many of today's e-commerce platforms operate under a *dual mode*: they not only provide online marketplaces where third-party sellers compete and offer products, but also act as retailers selling products of their own. We study how dual-mode operation changes price competition on e-commerce platforms. Using a stylized model that incorporates some of the key features of Amazon, we show that, if the platform both earns commissions on third-party sales and retails on its own marketplace, price competition changes along two margins: platform retail presence can raise or lower rivals' price levels, and common cost shocks need not be passed through symmetrically across seller types. Using data from Amazon U.S., we provide event-study evidence on third-party price dynamics around temporary changes in Amazon's retail availability. We find a 6% price decrease (5% price increase) in the days after Amazon enters (exits) as a reseller alongside third parties. Consistent with the model, we find asymmetric pass-through of commodity-cost shocks: third-party prices respond more strongly than Amazon's own retail prices, with pass-through approaching full transmission only at longer horizons.

Keywords: E-Commerce, Dual Mode Platform, Online Pricing, Pass-Through

JEL Classification: D4, L1, E3

*We thank Philippe Bacchetta, Luca Bennati, Brett Hollenbeck, and Morten Saethre for helpful comments, and Eduard Boehm, Clara Jean, and Marc Ivaldi for insightful discussions. We also thank participants at talks given at EARIE (2024), MaCCI (2025), the Paris Conference on Digital Economics (2026), the Bristol Empirical IO Workshop (2026), the Toulouse School of Economics, the Stockholm School of Economics, NHH Bergen, the Joint Research Centre of the European Commission, and Xiamen University for their questions and comments. We acknowledge generous funding through the HEC Lausanne Research Fund. Any errors are our own.

†Email: luise.eisfeld@unil.ch.
‡Email: yanjun@xmu.edu.cn.
§Email: yuli@lnu.edu.cn.

1 Introduction

A growing number of e-commerce platforms – from Amazon and Walmart to Zalando and JD.com – operate in *hybrid* mode: they host third-party sellers while also selling directly to consumers themselves. Amazon is the leading example. Roughly 60% of sales on the platform are generated by third-party sellers (Federal Trade Commission, 2024), from which Amazon earns ad valorem commissions; yet Amazon also retails many of the same products on its own account. This dual role gives the platform a distinctive stake in downstream competition: it benefits when third-party sellers make sales, but it also competes with them for the transaction itself. As online retail expands[1] and platforms increasingly combine marketplace intermediation with their own retail activity, the competitive implications of hybrid operation have become increasingly important.

In this paper, we study how this distinctive market structure shapes price competition in product markets. Dual-mode platforms are not well described by the familiar case of a retailer selling private-label products alongside branded goods. In that setting, the retailer typically controls the retail prices of all products it carries. On a dual-mode platform, by contrast, third-party sellers set their own prices. Nor is a dual-mode platform well captured by models of common ownership: through percentage fees, the platform takes a share of sellers' revenues rather than their profits, and thus does not internalize their costs. These features create pricing incentives that are specific to dual-mode platforms and absent from more familiar market structures.

To study how price competition plays out on hybrid platforms, we combine a theoretical model with reduced-form evidence from Amazon. On the theory side, we develop an oligopoly pricing model in which a dual-mode platform intermediates third-party sellers while also retailing on its own account. The model highlights two opposing forces of platform entry – a revenue-sharing effect, which softens price competition, and a commission-saving effect, which intensifies it – and implies that the platform should exhibit lower pass-through of common cost shocks than third-party sellers. On the empirical side, we collect a granular panel of product-level prices from `amazon.com` to document stylized facts on Amazon's participation as a first-party seller and to confront our theory predictions with the data. In short-run event studies around repeated within-ASIN changes in Amazon's retail availability, third-party prices decline by about 6% after Amazon appears as a seller and increase by about 5% after Amazon disappears. Moreover, our pass-through estimates show that third-party prices respond more strongly to commodity-cost shocks than Amazon prices, consistent with the model's prediction. Together, the entry/exit

[1] See, e.g., https://www.census.gov/retail/ecommerce.html for the United States, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=E-commerce_statistics_for_individuals for the European Union, https://www.ons.gov.uk/businessindustryandtrade/retailindustry/timeseries/ms6y/drsi for Great Britain, and https://dzswgf.mofcom.gov.cn/sjcx.html for China (accessed 28/08/2024).

and pass-through patterns show that dual-mode operation affects both third-party price levels and the incidence of upstream cost shocks. Pass-through is informative about this incidence, as it measures the extent to which cost shocks are transmitted to consumer prices, rather than absorbed in seller margins.[2]

Our oligopoly model formalizes this mechanism. A platform mediates price competition among differentiated-product retailers and, under a dual mode, also sells on its own account in direct competition with third-party (3P) sellers. We identify two platform-specific pricing wedges that arise when the platform sells its own products, relative to an otherwise comparable independent seller. First, because the platform earns an ad-valorem commission on third-party sales, raising its own price diverts demand toward third-party sellers and increases commission revenue. We refer to this as the revenue-sharing effect, which gives the platform an incentive to set a higher retail price. Second, because the platform does not pay the commission on its own sales, it faces a lower effective marginal cost than third-party sellers. We refer to this as the commission-saving effect, which gives the platform an incentive to set a lower retail price.

These two distinctive pricing incentives render the platform's entry effect different from that of an independent seller, and in general ambiguous. In particular, when the revenue-sharing effect is strong, the platform's entry may exert little competitive pressure on 3P sellers. As a result, equilibrium prices may decrease only mildly (or not at all), contrary to the common belief that platforms like Amazon inevitably price-squeeze 3P sellers through their own offerings.

However, even if the platform's entry does not immediately price-squeeze 3P sellers, the commission-saving effect places 3P sellers at a structural competitive disadvantage and makes them particularly vulnerable to cost shocks. We find that, under an industry-wide common cost shock, the platform's price responds less than that of a 3P seller. This arises because the commission fee inflates 3P sellers' effective marginal costs, forcing them to adjust their prices more aggressively in response to cost shocks. Our theoretical pass-through results suggest that, in episodes of positive cost shocks, 3P sellers therefore face a trade-off: either absorb the higher costs and compress their margins, or raise prices more sharply than the platform, which causes them to lose traffic and market share. Since cost shocks occur repeatedly in practice, these temporary disadvantages can accumulate over time, enabling the platform to strengthen its position and expand its market presence from a dynamic perspective.

Our empirical objective is to identify the pricing margins implied by the theory. Hence, we pair the model with reduced-form evidence, without imposing a structural model of demand and seller conduct. We use a granular panel of prices for tens of thousands of products in

[2]For a recent popular-press discussion of cost shocks faced by Amazon sellers, see https://www.reuters.com/business/retail-consumer/some-amazon-sellers-are-pulling-out-prime-day-amid-trump-tariffs-2025-04-28/ (accessed 20/11/2025).

narrowly defined product categories sold on amazon.com, the U.S.-facing storefront of Amazon. We first document new descriptive facts on the degree of first-party (Amazon) selling within and across narrow product categories. We find that Amazon enters as a re-seller very selectively, and that products sold by Amazon tend to be high-turnover products. Using hundreds of repeated within-ASIN episodes in which Amazon temporarily enters or exits as a retailer, we conduct short-run event studies of third-party pricing. Under the assumption that the timing of these entry and exit episodes is conditionally exogenous – an assumption we assess using pre-trends – Amazon re-entry is followed by a 6% decrease third-party prices, while Amazon exit is followed by a 5% increase. Our theory predicts that hybrid operation also changes pricing along a second margin, namely how prices react to common cost shocks. We estimate the impact of recent shocks to commodity costs, in particular, to cocoa and coffee beans, on consumer prices, and find that Amazon passes on common cost shocks to a smaller degree than third-party sellers. These empirical patterns line up with the commission-saving and revenue-sharing mechanisms in our model.

Related Literature

First, our paper contributes to the growing literature studying the implications of hybrid platforms for competition, innovation, and consumer welfare.[3] This literature has reached contrasting conclusions on whether allowing platforms to operate in dual mode harms consumers. Hagiu, Teh and Wright (2022) show that dual mode can benefit both the platform and consumers: the platform's own product "price-squeezes" a superior third-party seller, intensifying competition and thus increasing total transactions on the marketplace. In contrast, Anderson and Bedre Defolie (2024) find that the hybrid mode harms consumers because the platform optimally raises its commission fee relative to the pure marketplace mode in order to steer demand toward its own offerings. The higher fee reduces third-party participation and product variety, ultimately resulting in higher final prices.

Our paper differs from this literature in perspective and methodology. First, rather than endogenizing the platform's business-model choice, seller participation, or the commission fee, we take these elements as exogenous. This allows us to highlight the hybrid platform as a distinctive market structure and to examine how it shapes price competition. Second, by modeling marketplace competition as an oligopoly, we can capture rich strategic interactions between the

[3]This literature is partly motivated by the increasing prevalence of hybrid platforms across the economy as well as antitrust concerns associated with this business model; for example, the European Commission's cases on Amazon Marketplace and Amazon Buybox (cases AT.40462 and T.40703, see https://ec.europa.eu/competition/antitrust/cases1/202310/AT_40703_8990760_1533_5.pdf, or the U.S. FTC case, see https://www.ftc.gov/news-events/news/press-releases/2023/09/ftc-sues-amazon-illegally-maintaining-monopoly-power (both accessed 25/11/2025).

platform's and third-party sellers' prices. This strategic interdependence – combined with the platform's ability to share in its rivals' revenues through ad-valorem commissions – gives rise to the two key driving forces in our analysis: the revenue-sharing effect and the commission-saving effect. These incentives are absent in the earlier models: the homogeneous-preferences and fixed-fee framework of Hagiu et al. (2022) and the monopolistic-competition setting of Anderson and Bedre Defolie (2024) do not generate such pricing interactions. Finally, our paper offers a novel result by showing how the hybrid mode creates asymmetric pass-through of cost shocks between the platform and third-party sellers, and how this asymmetry can cumulatively disadvantage third-party sellers in dynamic competition. To the best of our knowledge, this channel has not been examined in the existing literature on hybrid platforms.[4]

Empirical evidence on pricing on Amazon remains relatively limited, especially on how hybrid operation shapes price competition and the transmission of cost shocks. Closest to our paper, a working paper by Crawford, Courthoud, Seibel and Zuzek (2022) use proprietary data from Amazon Germany's Home & Kitchen department to study Amazon's first-party entry and show that such entry is associated with lower third-party prices and modest market expansion. Using the Toys'R'Us bankruptcy as a shock to offline competition, He, Reimers and Shiller (2022) find that Amazon raises prices, consistent with competition from brick-and-mortar retailers. Other empirical work studies which products Amazon chooses to enter (Zhu and Liu, 2018), the existence and implications of self-preferencing and steering (Lee and Musolff, 2025; Chen and Tsai, 2024; Dendorfer and Seibel, 2026; Farronato, Fradkin and MacKay, 2023, 2025), the behavioral consequences of transparency about the Buy Box (Boldrini and Braulin, 2025), and the consequences of regulating Amazon. In particular, Gutiérrez (2022) studies the welfare consequences of regulating Amazon, while Bennati (2026) develops a structural model of Amazon Marketplace competition to quantify the price, entry, and welfare effects of restricting Amazon's own retail activity. Relative to this literature, our paper studies two pricing margins of competition on Amazon: the short-run effect of Amazon's retail presence on third-party prices, and the pass-through of common cost shocks across seller types. We thereby naturally relate to research that empirically estimates pass-through in various real-world settings (e.g., Pless and van Benthem, 2019; Fabra and Reguant, 2014; Miller, Osborne and Sheu, 2017; Ganapati, Shapiro and Walker,

[4]There exists further theoretical literature that deals with hybrid platforms, but does not address the pricing decisions that we study. For instance, this literature analyzes a hybrid platform's choice of entering into product markets (Etro, 2021); the optimal commission rate (e.g., Etro (2023)), or the optimal response of such a platform to collusive agreements among its sellers (Bisceglia and Padilla, 2023). Further theoretical work examines the incentives of platforms to "steer" consumers to their own products or to give biased recommendations (e.g., de Cornière and Taylor (2014); Drugov and Jeon (2019)), and estimates the existence of such steering empirically (see Lee and Musolff (2025) for Amazon, and Cure, Hunold, Kesler, Laitenberger and Larrieu (2022) in the context of a travel meta-search platform). Madsen and Vellodi (2023) study to what extent restrictions on a platform's use of data about the demand third-party sellers are facing, and the platform's subsequent entry decisions, can affect ex-ante product innovation incentives.

2020; Genakos and Pagliero, 2022; Montag, Mamrak, Sagimuldina and Schnitzer, 2023; Alvarez, Cavallo, MacKay and Mengano, 2024). Existing evidence of pass-through in e-commerce settings is macroeconomic, but offers opposite conclusions and does not model the competitive structure or pricing incentives of different seller types on e-commerce platforms (Cavallo, 2019; Gorodnichenko and Talavera, 2017; Goolsbee and Klenow, 2018).

Outline. This paper is structured as follows. Section 2 describes the main institutional features of e-commerce platforms like Amazon, introduces the data we have collected, and provides stylized facts on seller participation and prices. Section 3 proposes a stylized model of price competition on a hybrid platform. Section 4 shows results on the prevailing price level and pass-through on Amazon. Section 5 covers future work we intend to do. Section 6 concludes.

2 Institutional Setting and Stylized Facts

2.1 Hybrid Selling on Amazon Marketplace

Large e-commerce platforms increasingly combine a traditional retailing and a marketplace role in a hybrid structure: they operate a marketplace for third-party sellers and effectively tax these via a commission fee, while also selling first-party products in direct competition with those sellers (see Figure 1). The evolution of hybrid marketplaces has occurred from both directions: pure marketplaces (e.g., eBay, Zalando, and Taobao) integrating into retailing, possibly to better control quality, inventory, and data, and traditional retailers (such as Walmart, Zalando, or Fnac) opening their digital infrastructure to third-party sellers.[5] Hybrid operation may expand product variety and amplify buyer-seller network effects, but it also creates novel competitive incentives, as the platform compares full first-party retail margins with third-party commission revenues.

Amazon – the world's largest e-commerce platform – has operated in dual mode since 2000. Third-party sellers account for over 60% of sales on Amazon, and 560,000 sellers – predominantly small and medium-sized firms – were active on the U.S. marketplace in 2021 (Federal Trade Commission, 2024).[6] Amazon levies an *ad valorem* commission ("referral fee") on each mediated third-party transaction. The referral-fee schedule is public, varies by product category between about 5 and 20% (and, in some categories, may vary price tier), and changes infre-

[5]Examples of pure marketplaces integrating into retailing include eBay (see https://www.yo-kart.com/blog/how-does-ebay-make-money/) and Taobao (see https://www.digitalwebsolutions.com/blog/alibaba-business-model/). Examples of traditional retailers opening their digital infrastructure to third-party sellers include Walmart (see `https://tinyurl.com/2pyvv3eh`), Zalando (see https://www.brandongroup.it/en/zalando-partner-program-what-it-is-and-how-to-become-a-zalando-partner/) and Fnac (see `https://www.lengow.com/integrations/fnac/`; all sources last accessed 01/03/2026)).

[6]Amazon sellers made an average yearly revenue of $250,000 in 2023; see `https://sell.amazon.com/blog/amazon-stats` (accessed 27/09/2024).

Figure 1: Illustration of a hybrid platform

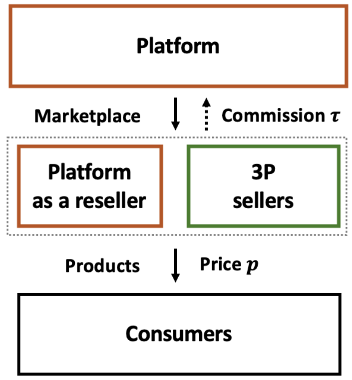

quently.[7] Amazon's first-party retail activity includes some private-label offerings (e.g., Amazon Basics), but most first-party sales are in fact of branded products. In the United States, Amazon likely has a large and loyal customer base, given the 190 million paid subscribers who benefit from free shipping in other perks.[8]

Unlike platforms like eBay, Amazon enforces a "single detail page" policy: rather than displaying fragmented search results where multiple sellers list the exact same item separately, Amazon aggregates all offers for a unique product, identified by an Amazon Standard Identification Number (ASIN), onto a single product page. This architecture impacts the nature of competition. Because sellers cannot differentiate themselves through separate listings, they compete directly for visibility within the shared page. The focal point of this competition is the "Buy Box" – the prominent section on the top-right of the page. Amazon uses an algorithmic mechanism to assign the Buy Box to a single seller (the "winner"), who captures an estimated 83% of sales for that product (Breslin, 2025). While the exact algorithm is proprietary, empirical evidence has shown that price and fulfillment method are primary determinants of winning the Buy Box (Chen, Mislove and Wilson, 2016; Lee and Musolff, 2025). Offers from non-winning sellers are relegated to a less visible list (often labeled "Other Sellers on Amazon") below the Buy Box. Third-party sellers can choose between managing their own logistics ("Fulfilled by Merchant", FBM) or outsourcing storage and shipping to Amazon ("Fulfilled by Amazon", FBA). Sellers selling products under FBA face, in addition to referral fees, FBA fulfillment fees for storage and delivery of the product.

[7] See https://sellercentral.amazon.com/help/hub/reference/external/GTG4BAWSY39Z98Z3 (accessed 27/11/2025); the most recent fee change on amazon.com we are aware of occurred in January 2024.

[8] See https://cirpamazon.substack.com/p/updated-us-amazon-prime-member-estimate (accessed 11/02/2026).

2.2 Data and Measurement

Building on the conceptual definition of a hybrid platform, we examine where direct competition between Amazon and third-party sellers arises, how prevalent it is, and how it varies across products and over time. We organize the data within narrowly defined Amazon product subcategories and treat the ASIN as the relevant product. We construct an ASIN-week panel that classifies each observation according to whether the ASIN is sold by Amazon, by third-party sellers, or by both. This panel forms the basis for the stylized facts and the empirical analyses that follow.

Data source. We collect data on prices of products sold on Amazon from Keepa (`keepa.com`), a company that provides data on prices, reviews, sales rank, and further both time-varying as well as cross-sectional attributes of products sold on Amazon by repeatedly scraping its product pages. Keepa updates its information multiple times per day, and its data have been used in previous academic literature (e.g., He et al. (2022); Gutiérrez (2022); Bennati (2026); Cabral and Xu (2021); He, Hollenbeck and Proserpio (2022)). Our sample period runs from January 2021 to March 2025, thereby excluding the period in which recent U.S. tariff changes may have influenced prices.

Selection of sample. Our analysis is conducted at the level of a subcategory. Because our focus is on pass-through, we study distinct subcategories that have experienced marginal-cost shocks in recent years. In particular, we examine products in the ground coffee and cocoa categories (see Section 4).[9] Grocery products now account for a substantial share of Amazon's revenue, suggesting that these categories capture mainstream products.[10] We also collect data for other food subcategories (wheat pasta) and non-food subcategories (TVs, washers, T-shirts, USB hubs, and SD cards), which we use as benchmarks. Within each product category, we restrict attention to the 10,000 products with the highest turnover, as proxied by both the current and the 180-day average Sales Rank. Overall, our sampling strategy lies between platform-wide studies that cover many departments with relatively few products per narrow category (e.g., Chen and Tsai, 2024; Crawford et al., 2022; Lee and Musolff, 2025) and single-market studies such as Bennati (2026).

[9] We do not study other products affected by cocoa-bean prices, such as chocolate bars, because these products often involve a large share of so-called "local costs" (Nakamura and Zerom, 2010), such as advertising expenditures. In addition, chocolate manufacturers have reportedly adjusted product recipes in response to input-price shocks (see `https://www.ft.com/content/169c9fca-1a7a-4c0f-84e9-398777335629`, last accessed 26/02/2025), which further limits the suitability of these products for studying price pass-through.

[10] In 2024, Amazon CFO Brian Olsavsky noted "strong growth in items like everyday essentials. This includes items like health, beauty, and personal care, as well as nonperishable grocery," and stated that "the strength in everyday essential's revenue is a positive indicator that customers are turning to us for more of their daily needs." See `https://www.fool.com/earnings/call-transcripts/2024/10/31/amazoncom-amzn-q3-2024-earnings-call-transcript/` (last accessed 01/02/2025).

In the Keepa data, whenever a product is being offered on Amazon, we observe the lowest third-party price fulfilled by Amazon (FBA), the lowest third-party price fulfilled by merchants (FBM), the lowest marketplace price overall (including Amazon retail), Amazon's own price, the total number of active sellers, the price displayed in the Buy Box, and the seller ID of the Buy Box winner. We moreover observe each product's sales rank ("Best Sellers Rank") over time, which aggregates recent and historical sales relative to the category peer group [11].

For the food subcategories, we recover package sizes (weight in grams) from product titles using the OpenAI API. Appendix A.2 provides details, and also shows that food products are sold predominantly in standard household sizes rather than bulk formats (Table 6).

We construct a product-day panel using prices recorded at midnight and then aggregate these data to the product-week level by averaging the price variables. Additional details on category definitions and the data collection procedure are provided in Appendix A.2.

Variable definitions. To capture the potentially strong cross-product heterogeneity in demand, we construct several variables that identify relatively high-turnover products. One coarse measure is an indicator for whether a product has any user reviews at the time of data collection. Because Amazon's sales-rank series occasionally exhibits jumps or breaks due to changes in category definitions, we also use two additional measures based on relative rank within category: (1) We classify an ASIN as *popular* if it falls in the top half of the contemporaneous sales-rank distribution within its category in more than 50% of observed weeks. We interpret this measure as capturing products that are persistently highly ranked, and hence relatively high-turnover, within their category. (2) For each product-week, we compute the product's *percentile in the sales-rank distribution* within its category. This variable equals 1 for the best-selling product and 0 for the worst-selling product in a given category-week.

Table 1 reports category-level summary statistics on sample coverage and seller participation. Because we query the data repeatedly over time, the final number of products in a category can substantially exceed 10,000. Product variety is substantial in every category: for example, the coffee category alone contains 17,250 distinct ASINs. Competitive conditions also vary meaningfully across categories. The average number of sellers per ASIN-week ranges from 1.3 for SD cards to 2.6 for coffee, and the share of ASIN-weeks with at least two observed sellers ranges from 15.0% to 41.9%. Review prevalence also differs markedly across categories. By contrast, Amazon private-label products account for only a very small share of products in every category, never exceeding 0.6% of the products scraped.

[11] See `https://sell.amazon.com/blog/amazon-best-sellers-rank` for further information (accessed 02/09/2024).

Table 1: Sample summary by category (ASIN-week panel; 222 weeks)

Category	Products (ASINs)	Avg. seller count (per ASIN-week)	ASIN-weeks with ≥ 2 sellers (%)	With reviews (%)	Amazon private label (%)
Cocoa	1,849	2.3	31.9%	64.8%	0.2%
Coffee	17,250	2.6	36.2%	88.8%	0.4%
Pasta	8,908	2.5	41.9%	71.8%	0.6%
Washers	3,652	1.7	22.0%	18.1%	0.0%
TV	7,237	2.4	32.6%	64.4%	0.0%
Shirts	25,989	1.7	17.4%	92.5%	0.0%
USB hubs	10,775	2.3	20.3%	34.4%	0.3%
SD cards	15,361	1.3	15.0%	14.0%	0.0%

Notes: We track sampled ASINs daily from January 1, 2021 to March 31, 2025 and collapse the data to the ASIN-week level. The average seller count corresponds to the average number of sellers listing an ASIN in a week, computed over ASIN-weeks with at least one observed seller, and includes Amazon retail when present. "With reviews" and "Amazon private label" are product-level shares. The review indicator is measured from the latest observed product page and is therefore time-invariant within the sample.

2.3 Stylized Facts

To provide context on how Amazon operates as a hybrid platform, we document four stylized facts covering two dimensions: Amazon's retail presence across categories and products, and the price levels and adjustment patterns of Amazon and third-party sellers.

Fact 1: Prevalence of dual mode across categories. *The prevalence of hybrid selling varies substantially across product categories, indicating that Amazon's direct retail participation is category-specific rather than a uniform platform-wide strategy.*

While over 60% of units sold on Amazon are supplied by third-party sellers (Federal Trade Commission, 2024), it is unclear that within a given product category, to what extent Amazon competes face-to-face with third-party sellers, and how Amazon's participation varies across categories. Figure 2 displays, for each category, the share of ASIN-week observations that are supplied exclusively by third-party sellers (3P-only), exclusively by Amazon (Amazon-only), or by both (dual-seller). Panel (a) weights each ASIN-week equally and shows that dual-seller states account for between about 0.8% and 8.1% of product-weeks across categories. Third-party sellers are the exclusive suppliers in at least 73.5% of product-weeks in every category, while Amazon is the exclusive seller in between 0.4% and 18.6% of product-weeks. Overall, Amazon's retail participation varies markedly across product categories: Amazon is active (either alone or alongside 3P sellers) in roughly 26% of ASIN-weeks in *Coffee* and about 24% in *TVs*, but in less than 2% of ASIN-weeks in *SD Cards*.

Fact 2: Presence of Amazon across products. *Amazon's direct retail participation is concentrated in more popular, higher-turnover products, while the vast majority of lower-turnover ASINs are supplied*

Figure 2: Extent of hybrid selling across categories

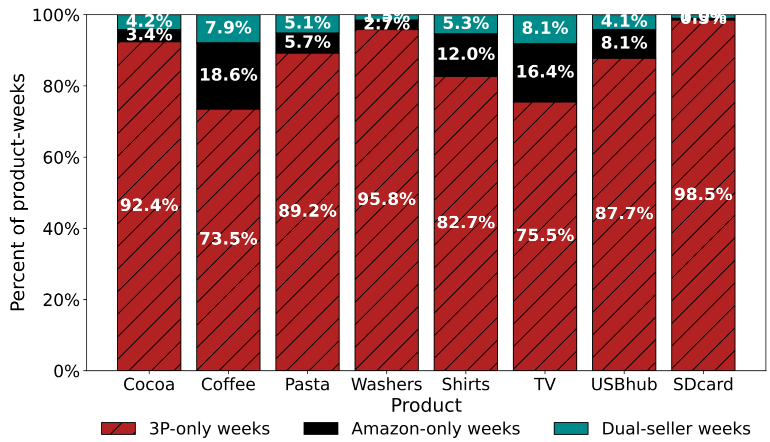

(a) Equal-weighted product-weeks.

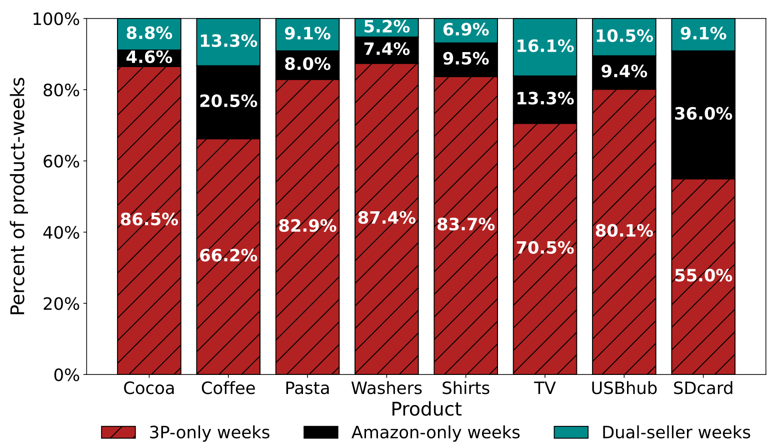

(b) Weighted by sales-rank percentile.

Notes: Each bar shows the share of ASIN-week observations that are 3P-only, Amazon-only, or dual-seller for a given category. Panel (a) weights each ASIN-week equally. Panel (b) weights ASIN-weeks by within-category sales-rank percentile, giving more weight to higher-turnover products. The sample covers ASINs observed weekly from January 2021 through March 2025. A product-week is classified as 3P-only (Amazon-only) if at least one 3P (Amazon) offer is recorded and the other side is absent; it is classified as dual-seller if both offer the products in the same week.

exclusively by third-party sellers.

Panel (b) of Figure 2 weights ASIN-weeks by their within-category sales-rank percentile, giving more weight to higher-turnover products. Comparing the equal-weighted shares in panel (a) and the sales-rank-weighted shares in panel (b) of Figure 2 shows that the prevalence of Amazon-only and dual-seller ASIN-weeks increases for all categories except for *Shirts*. For instance, in both *TVs* and *Coffee*, the share of ASIN-weeks in which Amazon is active (Amazon-only or dual-seller) rises from about 24-26% under equal weighting to around 34-30% under sales-rank percentile weighting. For *SD Cards*, Amazon-only and dual-seller states account for around 36%

and 9% of sales-rank-weighted ASIN-weeks, respectively, which is markedly larger than before. Taken together, these patterns indicate that Amazon's direct retail presence is relatively more common among higher-turnover ASINs, while third-party sellers appear to account for most of the low-turnover "long tail" within categories[12]. Figure 8 in Appendix A.4 replicates Figure 2 panel (a) for the subset of popular products, as defined in Section 2.2, with qualitatively similar findings.

Figure 3: Percent of products for which Amazon prices higher than the lowest third-party

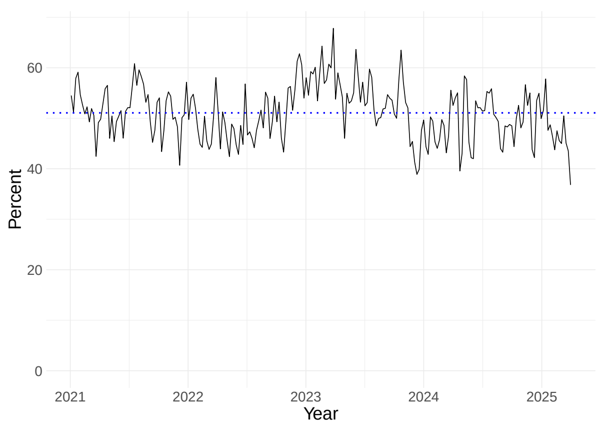

(a) Cocoa.

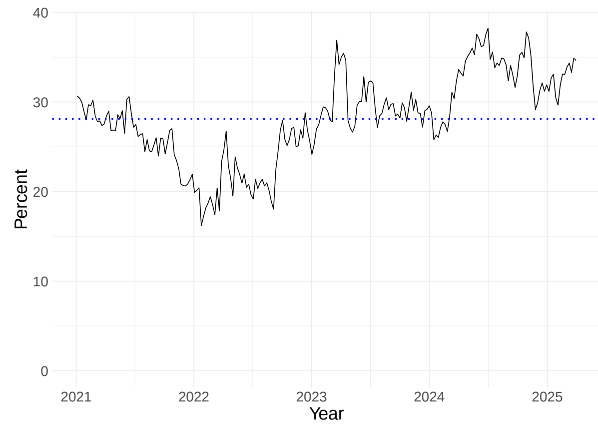

(b) Coffee.

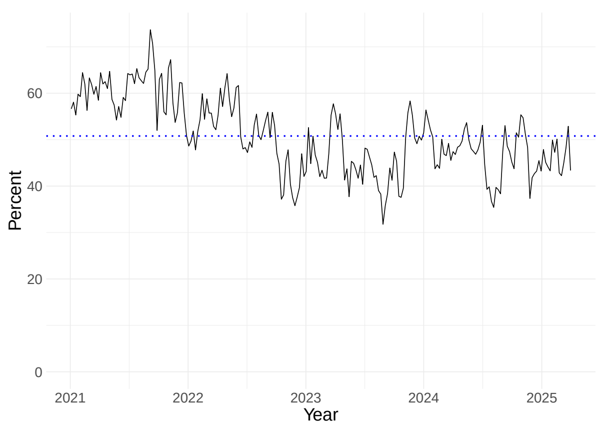

(c) SD Cards.

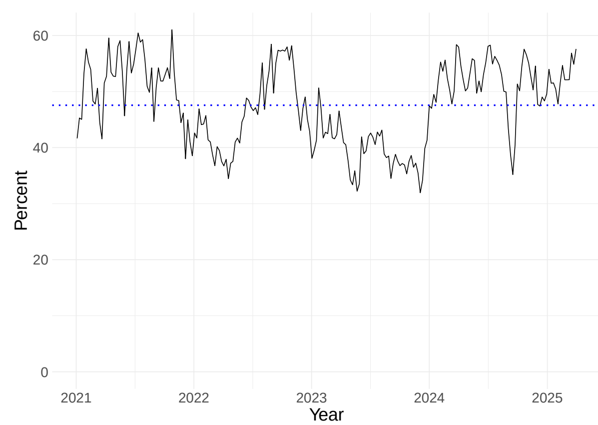

(d) USB Hubs.

Notes: Plots use product-week observations in which both a third-party seller's, as well as Amazon's price is available (i.e., both seller types are active). Blue dotted line indicates median.

Fact 3: Price levels of Amazon and 3P products. *In the product categories in our sample,*

[12]These patterns should not be given a causal interpretation, as from the figures alone, it is unclear whether Amazon chooses to start selling popular ASINs, or whether Amazon's retail presence renders ASINs more popular. Nevertheless, these patterns are consistent with prior work documenting that Amazon tends to enter more popular product niches and that third-party sellers disproportionately populate the long tail (Bennati, 2026; Crawford et al., 2022; Zhu and Liu, 2018). Appendix A.6 summarizes what researchers have learned about the determinants of Amazon's presence as a retailer, and reports the corresponding correlates of Amazon's presence in our own data.

conditional on identical products being simultaneously offered by both Amazon and third-party sellers, Amazon does not systematically charge higher prices than third-party sellers.

Since Amazon disproportionately sells more popular products, a raw comparison of Amazon and third-party prices across all ASINs would conflate pricing behavior with product selection. We therefore restrict attention to ASIN-week observations in which both Amazon and at least one third-party seller simultaneously offer the same ASIN. Figure 3 plots, for each product category, the fraction of product-week observations in which Amazon's posted price is strictly higher than that of the lowest-priced third-party seller for the same ASIN. Across all categories, this fraction fluctuates around 30-50%, indicating that Amazon is not predominantly the high-price seller when identical products are directly comparable. This descriptive fact is confirmed after controlling for ASIN and year-week fixed-effects, and accords with Bennati (2026)'s descriptive evidence. Even when we consider different subsamples, our regression estimates confirm that Amazon does not charge systematic price premia relative to third-party sellers, although the magnitude of the relative log price difference between Amazon and third-party sellers varies somewhat across categories (see Appendix A.5).

Fact 4: Price adjustments of Amazon and third-party products. *In the product categories in our sample, conditional on identical products being simultaneously offered by both Amazon and third-party sellers, Amazon does not systematically adjust its products' prices at a higher frequency than third-party sellers.*

To assess whether this pattern holds beyond the descriptive level, we follow the same sample restriction as in Fact 3 and construct a stacked dataset in which each ASIN-week observation contributes two rows — one for Amazon and one for the third-party seller — allowing us to compare adjustment frequency within the same product and time period. We define a price adjustment as a week-to-week log price change exceeding a given threshold (0.5% and 1%), and regress an indicator for adjustment on an Amazon dummy, controlling for ASIN fixed effects and time fixed effects. Table 2 reports the results. Across all specifications and both thresholds, the coefficient on the Amazon indicator is small in magnitude and statistically insignificant, indicating that Amazon does not adjust prices at a systematically higher frequency than third-party sellers on jointly-offered products. We further confirm this result on a broader sample that relaxes the strict simultaneity requirement, with consistent findings across specifications.

In summary, the data reveal a market structure in which Amazon's direct retail presence is heterogeneous across categories and concentrated in more popular products. On the products where both seller types compete directly, prices and adjustment frequencies are broadly similar except for coffee where Amazon sets lower prices, suggesting that head-to-head competition shapes pricing behavior in ways that compress differences between the platform's retail arm

Table 2: Amazon vs. Third-Party Price Adjustment Frequency

Dependent Variables:	is_adjust_05	is_adjust_10	is_adjust_05	is_adjust_10
Model:	(1)	(2)	(3)	(4)
Amazon indicator	0.0162	0.0070	0.0210	0.0085
	(0.0328)	(0.0308)	(0.0318)	(0.0292)
ASIN	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes
Observations	12,612	12,612	163,513	163,513
R^2	0.13797	0.12799	0.25547	0.22185

Notes: Dependent variable equals one if the absolute weekly log price change exceeds the threshold shown in each column (0.5% and 1% respectively), and zero otherwise. Joint sample restricts to ASIN-week observations with non-missing Amazon and third-party prices in both t and $t-1$; broad sample uses all seller-week observations with non-missing own prices in t and $t-1$. All columns include ASIN and week fixed effects. Clustered (ASIN) standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$.

and third-party sellers. However, from the static descriptive statistics alone, we cannot observe how such features are formed and sustained dynamically. For example, when Amazon enters a product previously sold exclusively by third-party sellers, does the observed price similarity reflect Amazon responding to existing third-party prices, or third-party sellers adjusting in response to Amazon's entry? Understanding the direction of this adjustment helps us interpret the price similarity documented in Fact 3. Furthermore, when adjustment frequencies appear similar across seller types, do Amazon and third-party sellers transmit upstream cost shocks – in terms of magnitude – to consumers to the same degree? The remainder of the paper addresses these two questions.

3 Theory

To interpret the forces driving pricing patterns and to decompose the mechanisms behind them, we next develop a theoretical framework that models the platform's dual objective within a hybrid structure.

3.1 General Framework

Suppose that a monopoly e-commerce platform P operates a marketplace in a given product market. There are n retailers, each selling a differentiated product on the marketplace. The demand faced by retailer $i \in \{1, \ldots, n\}$, denoted $D_i(\mathbf{p}) = D_i(p_i, \mathbf{p}_{-i}) = D_i(p_1, \ldots, p_n)$, depends on its own price p_i and the prices of its rivals $\mathbf{p}_{-i}$. We assume a symmetric demand system. The marginal cost of retailer $i \in \{1, \ldots, n\}$ is c_i. The platform charges a commission fee τ on the revenue generated by each third-party (3P) seller on the platform. We take the commission fee as exogenous.

We consider two possible business models for the platform: the pure marketplace mode and the hybrid mode. In the pure marketplace mode, the platform hosts only third-party products. In the hybrid mode, the platform not only runs the marketplace but also sells its own product and competes directly with 3P sellers.

Timing: Under each business model of the platform, all retailers simultaneously choose their prices $p_i, i \in \{1, 2, \ldots, n\}$.

3.2 Pricing Incentives

In this subsection, we examine how the hybrid mode—that is, the platform's entry as a retailer—affects price competition in the product market, and in particular, how it influences the prices of third-party (3P) sellers relative to the pure marketplace mode. To isolate the pure effect of the platform's role as a retailer on pricing incentives from the entry effect, we keep the number of retailers fixed at n by assuming that retailer 1 represents the platform under the hybrid mode, and we denote its price by p_P.

In the pure marketplace mode, each retailer $i \in \{1, 2, \ldots, n\}$ chooses its price p_i to solve

$$\begin{aligned} &\max_{p_i} \, [(1-\tau)p_i - c_i]D_i(\mathbf{p}) \\ \Leftrightarrow &\max_{p_i} \, (1-\tau)(p_i - \frac{c}{1-\tau})D_i(\mathbf{p}) \end{aligned}$$

The first-order condition is

$$D_i(\mathbf{p}) + (p_i - \frac{c_i}{1-\tau})\frac{\partial D_i(\mathbf{p})}{\partial p_i} = 0, \quad i = 1, \ldots, n \tag{1}$$

Therefore, under the pure marketplace mode, the commission fee effectively scales up each retailer's marginal cost by a factor of $\frac{1}{1-\tau}$. We denote the equilibrium under the pure marketplace mode by $\mathbf{p}^* = (p_1^*, \ldots, p_n^*)$, characterized by Equation (1).

In the hybrid platform mode, 3P retailers $i = 2, \ldots, n$ face the same profit maximization problem as under the pure marketplace mode. In contrast, the platform as retailer 1 solves the following problem

$$\max_{p_P} \, (p_P - c_P)D_P(\mathbf{p}) + \sum_{i=2}^{n} \tau p_i D_i(\mathbf{p}),$$

where its marginal cost is $c_P = c_1$.

The first-order condition is

$$D_P(\mathbf{p}) + (p_P - c_P)\frac{\partial D_P(\mathbf{p})}{\partial p_P} + \tau \sum_{i=2}^{n} p_i \frac{\partial D_i(\mathbf{p})}{\partial p_P} = 0 \tag{2}$$

We denote the equilibrium under the hybrid mode by $\hat{\mathbf{p}}^* = (\hat{p}_P^*, \ldots, \hat{p}_n^*)$.

Proposition 1. *When the platform operates in the hybrid mode, its pricing incentive as a retailer differs from that of an independent 3P seller in two respects:*

(1) ***Revenue-sharing effect****. Through the term $\tau \sum_{i=2}^{n} p_i \frac{\partial D_i(\mathbf{p})}{\partial p_P}$, the platform internalizes the positive impact of raising its own price on the sales of 3P sellers, and hence on its commission revenue. This gives it an incentive to set a higher price than p_1^*.*

(2) ***Commission-saving effect****. Because the platform does not pay a commission fee to itself, its effective marginal cost is lower than that of 3P sellers. This cost advantage provides an incentive to set a lower retail price than p_1^*.*

To facilitate comparative statics, we now impose the following specification.

Assumption 1. *All firms have identical marginal costs $c_i = c, \forall i \in \{1, \ldots, n\}$ and symmetric demand functions:*

$$D_i(\mathbf{p}) = \frac{1 - \gamma - (1 + (n-2)\gamma)\, p_i + \gamma \sum_{j \neq i} p_j}{(1-\gamma)\,(1 + (n-1)\gamma)},\ \forall i \neq j \in \{1 \ldots, n\}.$$

Under this assumption, we focus on the symmetric equilibrium under the pure marketplace mode. We denote the equilibrium price with n competing retailers by $p_M^*(n)$. Similarly, under the hybrid mode, we focus on the semi-symmetric equilibrium in which all 3P sellers charge the same price. We denote the platform's price by $\hat{p}_P^*(n)$ and the common price charged by 3P sellers by $\hat{p}_S^*(n)$.

Proposition 2. *Under Assumption 1, when $0 < \tau \leq \max\left\{0, 1 - \frac{2c(1+(n-2)\gamma)}{(n-1)\gamma}\right\}$, the revenue-sharing effect dominates the commission-saving effect. Consequently,*

(i) the platform charges a higher equilibrium price than 3P sellers: $\hat{p}_P^*(n) > \hat{p}_S^*(n)$;

(ii) the platform's entry exerts a smaller downward pressure on 3P sellers' prices than entry by an otherwise identical 3P seller: $\hat{p}_S^*(n) - p_M^*(n-1) > p_M^*(n) - p_M^*(n-1)$.

At first glance, this result may appear counterintuitive, as the revenue-sharing effect dominates when the commission rate τ is relatively low. In fact, this outcome is driven by the fact that the 3P sellers' effective marginal cost, $\frac{c}{1-\tau}$, increases nonlinearly with the commission rate. As a result, the commission-saving effect becomes increasingly important as τ rises. Only when the commission rate is sufficiently high does the commission-saving effect become strong enough to outweigh the revenue-sharing effect, inducing the platform to charge a lower price than 3P sellers.

Moreover, whenever the revenue-sharing effect dominates, the platform's entry exerts a smaller downward pressure on 3P sellers' prices than entry by an otherwise identical 3P seller. In other words, despite making the marketplace more competitive, the platform does not necessarily engage in aggressive price competition with 3P sellers by exploiting its cost advantage. This is precisely because the platform internalizes the positive impact of higher 3P sales on its commission revenue.

The intuition behind why the conditions for these two results coincide is straightforward. Under the linear demand specification, prices are strategic complements. Suppose the revenue-sharing effect dominates the commission-saving effect when evaluating the platform's marginal profit from increasing p_P at the pure-marketplace equilibrium $\mathbf{p}_M^*(n)$ (i.e., the left-hand side of Equation (2) is positive). In that case, the platform finds it profitable to raise its price above $p_M^*(n)$. Because prices are strategic complements, third-party sellers respond by raising their prices as well. Consequently, in equilibrium, we have $\hat{p}_P^*(n) > \hat{p}_S^*(n) > p_M^*(n)$ Conversely, if the commission-saving effect dominates, the platform finds it profitable to set a price below $p_M^*(n)$. Third-party sellers then respond by lowering their prices. As a result, the hybrid-mode equilibrium features lower prices for all retailers than the pure marketplace equilibrium.

3.3 Asymmetric Pass-Through

In this subsection, we further investigate how the platform's special pricing incentives under the hybrid mode affect pass-through – that is, how the platform and 3P sellers respond differently to cost shocks.

To make the analysis concrete, we focus on the case with $n = 2$ and denote the two sellers by P and S. In addition, we assume the following linear demand function:

$$D_i(p_i, p_j) = \frac{1 - \gamma - p_i + \gamma p_j}{1 - \gamma^2}, i, j \in \{P, S\}; i \neq j,$$

where $0 < \gamma < 1$ measures the degree of product homogeneity.

Substituting the demand functions into the first-order conditions in Equation (1) and Equation

(2), the equilibrium prices under the hybrid mode are given by:

$$p_P = \frac{2c_P(1-\tau) + \gamma(1+\tau)c_S + (1-\gamma)(1-\tau)[2+\gamma(1+\tau)]}{(1-\tau)(4-\gamma^2(1+\tau))},$$

$$p_S = \frac{2c_S + \gamma c_P(1-\tau) + (1-\gamma)(1-\tau)(2+\gamma)}{(1-\tau)(4-\gamma^2(1+\tau))}.$$

Proposition 3. Pass-through matrix of cost shocks. *The pass-through from the platform's and the third-party seller's cost shocks to equilibrium prices is given by*

$$\rho = \begin{bmatrix} \frac{\partial p_P}{\partial c_P} & \frac{\partial p_P}{\partial c_S} \\ \frac{\partial p_S}{\partial c_P} & \frac{\partial p_S}{\partial c_S} \end{bmatrix} = \frac{1}{4-\gamma^2(1+\tau)} \begin{bmatrix} 2 & \frac{\gamma(1+\tau)}{1-\tau} \\ \gamma & \frac{2}{1-\tau} \end{bmatrix}.$$

(1) The platform has a lower own-cost pass-through rate than the 3P seller: $\frac{\partial p_P}{\partial c_P} < \frac{\partial p_S}{\partial c_S}$.

(2) The platform has a higher cross-cost pass-through than the 3P seller: $\frac{\partial p_S}{\partial c_P} < \frac{\partial p_P}{\partial c_S}$.

(3) When there is an industry-wide cost shock ε to both sellers' costs, the platform experiences a smaller price increase than the 3P seller: $\frac{dp_P}{d\varepsilon} < \frac{dp_S}{d\varepsilon}$.

Proposition 3 illustrates the crucial difference between own-cost and cross-cost pass-through under the hybrid mode. To understand the intuition, we rewrite the equilibrium conditions as

$$\begin{cases} p_P - (c_P + \tau p_S A_P) = -\frac{D_P}{\partial D_P/\partial p_P} & \text{(Platform)} \\ p_S - \frac{c_S}{1-\tau} = -\frac{D_S}{\partial D_S/\partial p_S} & \text{(3P Seller)} \end{cases}$$

Here, $A_P = -\frac{\partial D_S(\mathbf{p}^*)/\partial p_P}{\partial D_P(\mathbf{p}^*)/\partial p_P} = \gamma$ is the diversion ratio from the platform's product to the 3P seller, measuring the fraction of platform's lost sales that can be captured by the 3P seller when the platform increases its price. Hence, the $\tau p_S A_P$ can be interpreted as the opportunity cost faced by the platform when selling an additional unit of its own product[13].

From the equilibrium conditions, the difference in own-cost pass-through rates arises from the commission-saving effect. The commission fee scales up the 3P seller's effective marginal cost, but not the platform's. This amplification—captured by the term $\frac{1}{1-\tau}$—makes the 3P seller inherently more sensitive to its own cost shocks than the platform. In contrast, the platform exhibits a higher cross-cost pass-through due to the revenue-sharing mechanism. When a positive cost shock to the 3P seller increases its price p_S, the platform's opportunity cost $\tau p_S A_P$ of selling

[13]The need to write pass-through in the form of a matrix does not rely on a particular form of demand function, but only on the asymmetric market structure imposed by the hybrid mode. We have also derived a general form of pass-through matrix under a hybrid platform à la Weyl and Fabinger (2013), but omitted it due to it being difficult to interpret.

its product rises. This gives the platform an additional incentive to increase its own price. As the 3P seller earns only sales revenue, it does not internalize a similar channel when the platform's cost rises.

Since own-cost effects are first-order and dominate cross-cost effects, the platform's aggregate pass-through rate under an industry-wide cost shock is smaller than that of the 3P seller.

In short, Proposition 3 shows that the platform's price is more stable in the face of cost shocks, whereas the 3P seller's price is more volatile and vulnerable. This asymmetry in price stability has broader competitive implications. In episodes of positive cost shocks, 3P sellers face greater pressure to raise their prices, causing them to lose traffic and market share to the platform's product. Since cost shocks occur repeatedly in practice, these temporary disadvantages can accumulate over time, allowing the platform to strengthen its position and expand its market presence from a dynamic perspective.

4 Empirical Evidence

Our theoretical framework highlights a pricing asymmetry between the platform-retailer and third-party sellers on a hybrid platform. Because the platform both earns commissions on third-party sales and does not pay commissions on its own sales, its pricing incentives differ from those of an otherwise similar third-party seller along two margins:

- **Prediction 1 (Effect on Equilibrium Price Level):** Compared to entry by an otherwise identical third-party seller, the platform's entry as a retailer puts greater downward pressure on incumbent third-party sellers' prices when the commission-saving effect dominates.

- **Prediction 2 (Asymmetric Pass-Through):** Due to the wedge created by commission fees, third-party sellers are more sensitive to cost shocks than the platform. Proposition 3 explicitly predicts that the platform's pass-through rate of common cost shocks should be strictly lower than that of third-party sellers ($\frac{dp_P}{d\varepsilon} < \frac{dp_S}{d\varepsilon}$).

In this section, we bring the model's predictions to the data. First, to study the price-level margin, we use repeated within-ASIN changes in Amazon's retail availability and estimate event-time patterns in third-party prices around these changes. Second, to speak to the asymmetric pass-through prediction, we use granular commodity cost shocks to compare the responsiveness of Amazon and third-party sellers to exogenous input price changes. Although the result that the platform has a lower pass-through of common cost shocks than third-party sellers is derived in a linear duopoly set-up, we view it as capturing a general mechanism of the model and use it to motivate our empirical analysis.

4.1 Amazon Retail Presence and Third-Party Price Levels

We first study the price-level margin of competition on a hybrid platform by examining how third-party prices evolve around changes in Amazon's retail availability. Cross-sectional comparisons are difficult to interpret because Amazon selectively sells particular products. We therefore restrict attention to ASINs that experience both Amazon and third-party participation at some point in the sample, and use short-run within-ASIN variation in whether an Amazon retail offer is observed. These repeated availability spells allow us to compare third-party prices for the same product before and after Amazon appears or disappears as a seller.

We broadly relate to reduced-form analyses of entry on equilibrium prices (for instance, Arcidiacono, Ellickson, Mela and Singleton (2020) for supermarkets, Fischer, Martin and Schmidt-Dengler (2025) for gasoline stations, and Crawford et al. (2022) for the Home & Kitchen segment on Amazon Germany), but differ in important ways. First, rather than exploiting a single entry event per ASIN with a never-treated control group, our within-ASIN study of repeated entry and exits avoids concerns about the comparability of never-treated ASINs. We use the terms "entry" and "exit" as shorthand for changes in observed Amazon retail availability within an ASIN, rather than necessarily permanent product-market entry or exit decisions. We can assess the validity of the identifying assumption using pre-event price trends. Second, by using both entry and exit events, we can examine whether the effects of Amazon's presence on third-party prices are reversible upon exit – a test that single-entry studies cannot perform.[14]

Amazon presence spells as identifying variation. Our identification strategy exploits within-ASIN variation in Amazon's retail presence, using the frequent and irregular spells during which Amazon enters or exits as a seller of a given ASIN. We restrict attention to ASINs that exhibit meaningful variation in seller composition: specifically, we require that (i) Amazon and third-party sellers are simultaneously active for at least 10 consecutive days, and (ii) only third-party sellers are active for at least 10 consecutive days.

Two features of the data are useful for assessing the plausibility of this design. First, there is substantial within-ASIN variation in Amazon retail availability: for ASINs that switch Amazon-presence status, Amazon-present and Amazon-absent spells recur frequently and vary considerably in duration (see Appendix A.7.1). Second, qualifying entry and exit events do not appear to be concentrated on a small number of calendar dates. This reduces concern that the estimates are driven solely by platform-wide shocks common to many ASINs, although it does not rule out ASIN-specific demand, inventory, or replenishment shocks.

[14] It should be noted that one should not interpret the adjustment speeed from these event studies, as Keepa may record changes in FBA or FBM prices only with a time lag.

Event definition. As Amazon-presence spells can be short and are not always preceded or followed by a stable period of third-party-only selling, we restrict attention to "clean" events with well-defined pre- and post-event windows. We define an entry event as a sequence in which (i) a third-party price is observed for at least K consecutive days and (ii) both an Amazon price and a third-party price are observed for the subsequent K consecutive days. Analogously, we define an exit event as a sequence in which (i) both Amazon and third-party prices are observed for at least K consecutive days and (ii) only a third-party price is observed for the subsequent K consecutive days. In the baseline specification, $K = 10$. The Appendix reports analogous results for $K = 15$.

Econometric specification. After isolating these types of events, we perform a sharp event study (Kleven, Landais and Søgaard, 2019) around symmetric windows of these events. Let

$$p_{it}^{3P,m} \equiv \text{lowest 3P offer price for ASIN } i \text{ on day } t \text{ under mode } m \in \{\text{FBA}, \text{FBM}\} \tag{3}$$

We write:

$$\log p_{it}^{3P,m} = \sum_{k=-10, k\neq -1}^{10} \beta_k^m \mathbf{1}\{\text{ET}_{it} = k\} + \alpha_i + \gamma_t + \varepsilon_{it}^m \tag{4}$$

where k indicates the event time, $\mathbf{1}\{\text{ET}_{it} = k\}$ is an indicator variable that is equal to 1 when the event time is equal to k, and α_i and γ_t are ASIN and calendar date fixed effects, respectively. We hence study the evolution of the lowest 3P offer price in a time window around qualifying event dates. The event study approach enables us to understand the dynamic adjustment of third-party seller's prices around the time of entry and exit. ASIN fixed effects control for both unobserved (e.g., date of product introduction) and observed (e.g., package size) characteristics that are time-invariant. Calendar date fixed effects can control for unobserved shocks to demand or supply that might affect the entire category. A causal interpretation requires as underlying assumptions that (1) the date of entry/exit is exogenous, and (2) that prices would evolve smoothly absent Amazon's entry or exit. We can visually assess the validity of these assumptions using the event study estimates.

We verify that our qualifying events are not clustered in time to rule out unobserved time shocks driving our effects (see Figure 11 in Appendix A.7.1 for visual evidence of qualifying events).

Results. Our key results are illustrated in Figure 4 in the case of FBA prices. Coefficients and fit statistics for both FBA and FBM prices are displayed in Table 3. For entry events (columns (1) and (2), we observe pre-event coefficients that are jointly insignificant, and a strong decline in the third-party seller's price by almost 6% in the case of the FBA price, and by 2% in the case of the FBM price. Looking at exit events (columns (3) and (4), we find the opposite pattern: mostly

Figure 4: Event study of entry and exit by Amazon as a seller into a given ASIN. Coffee data (specifications (1) and (3) of Table 3).

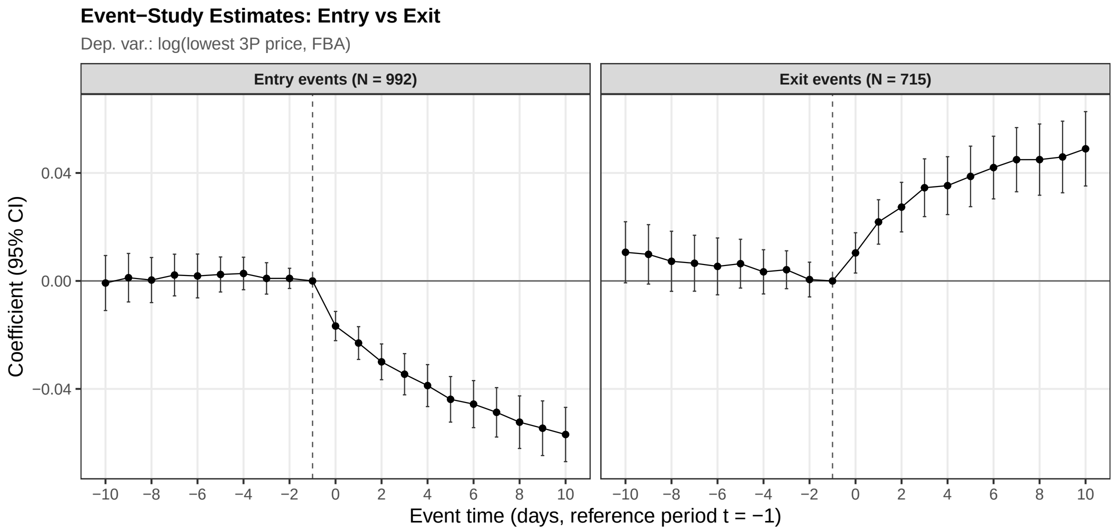

flat pre-trends are followed by an increase in the price by up to 5% in the case of FBA prices, and by 3% in the case of FBM prices.

Focusing separately on events at which Amazon enters with a relatively high (low) price yields intuitive results (see Appendix A.7.2): price declines upon entry are significant only if Amazon enters with a price that is lower than the average FBA price in the pre-window period; the analogous finding arises in the case of exit events. We also obtain very similar results when extending the event windows to $K = 15$. Finally, without being able to make a causal statement, we observe that the extent of price change differs slightly by the extent of overall third-party competition; for exit events in particular, prices tend to increase more post-exit in ASINs with a *higher* level of competition, although the pattern is somewhat noisy. Finally, the Appendix also shows the results from a more "naive" standard two-way fixed effects specification that we conduct using daily data across multiple product categories: we consistently find a negative and relation between Amazon's presence and the prevailing third-party seller's price.

Discussion and further robustness. By focusing our analysis on products which Amazon eventually enters in, we believe that we can control for a possible selection effect that stems from the fact that Amazon enters products selectively (as documented in Fact 2 above). However, positive demand shocks that occur within an ASIN over time might be correlated with both higher prices and Amazon entry, and this endogeneity could bias our coefficient β upwards (i.e., towards zero or positive). Further, Amazon exit might be a result of lack of demand, or a symptom driven by a stock-out that occurs after a surge in demand. These forces could bias our estimate in different directions. While our fairly flat pre-trends do not suggest such effects, we

Table 3: Amazon Entry and Exit Event-Studies: log(FBA) and log(FBM) Outcomes

Dependent Variables:	log(FBA price)	log(FBM price)	log(FBA price)	log(FBM price)
Specification	Entry, log(FBA)	Entry, log(FBM)	Exit, log(FBA)	Exit, log(FBM)
Model:	(1)	(2)	(3)	(4)
event_time = -10	-0.0008	0.0073*	0.0106*	0.0051
	(0.0052)	(0.0044)	(0.0058)	(0.0047)
event_time = -9	0.0012	0.0060	0.0098*	0.0065
	(0.0046)	(0.0039)	(0.0056)	(0.0049)
event_time = -8	0.0003	0.0071*	0.0073	0.0036
	(0.0043)	(0.0039)	(0.0057)	(0.0044)
event_time = -7	0.0022	0.0054	0.0066	0.0064
	(0.0039)	(0.0038)	(0.0053)	(0.0047)
event_time = -6	0.0018	0.0034	0.0054	0.0082*
	(0.0041)	(0.0035)	(0.0054)	(0.0048)
event_time = -5	0.0024	0.0047	0.0064	0.0002
	(0.0033)	(0.0034)	(0.0046)	(0.0039)
event_time = -4	0.0028	0.0019	0.0034	0.0015
	(0.0031)	(0.0033)	(0.0042)	(0.0037)
event_time = -3	0.0009	0.0022	0.0041	0.0040
	(0.0030)	(0.0026)	(0.0036)	(0.0037)
event_time = -2	0.0009	0.0006	0.0005	0.0036
	(0.0019)	(0.0019)	(0.0033)	(0.0029)
event_time = 0	-0.0167***	-0.0076***	0.0104***	0.0103***
	(0.0028)	(0.0022)	(0.0038)	(0.0028)
event_time = 1	-0.0230***	-0.0140***	0.0219***	0.0174***
	(0.0031)	(0.0032)	(0.0042)	(0.0037)
event_time = 2	-0.0300***	-0.0158***	0.0273***	0.0268***
	(0.0034)	(0.0035)	(0.0047)	(0.0045)
event_time = 3	-0.0346***	-0.0157***	0.0345***	0.0326***
	(0.0039)	(0.0038)	(0.0055)	(0.0043)
event_time = 4	-0.0388***	-0.0143***	0.0353***	0.0319***
	(0.0040)	(0.0039)	(0.0055)	(0.0049)
event_time = 5	-0.0439***	-0.0160***	0.0387***	0.0331***
	(0.0043)	(0.0039)	(0.0057)	(0.0048)
event_time = 6	-0.0456***	-0.0156***	0.0420***	0.0323***
	(0.0044)	(0.0044)	(0.0059)	(0.0049)
event_time = 7	-0.0487***	-0.0191***	0.0449***	0.0328***
	(0.0047)	(0.0041)	(0.0061)	(0.0051)
event_time = 8	-0.0523***	-0.0225***	0.0449***	0.0340***
	(0.0050)	(0.0043)	(0.0067)	(0.0054)
event_time = 9	-0.0546***	-0.0205***	0.0459***	0.0317***
	(0.0052)	(0.0045)	(0.0068)	(0.0053)
event_time = 10	-0.0569***	-0.0216***	0.0490***	0.0327***
	(0.0051)	(0.0047)	(0.0070)	(0.0057)
ASIN	Yes	Yes	Yes	Yes
Date	Yes	Yes	Yes	Yes
Observations	20,832	30,450	15,015	22,281
Number of Events	992	1,450	715	1,061
Number of ASINs	644	880	477	690
R^2	0.96986	0.96490	0.97178	0.96605

Notes: Two-way clustered (ASIN × date) standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$. Columns (1) and (2) report entry event studies; columns (3) and (4) report exit event studies. Dependent variables are log(FBA price) in columns (1) and (3) and log(FBM price) in columns (2) and (4). Coefficients are event-time indicators relative to $t = -1$ (omitted period). Entry columns: sample of all qualifying Amazon-entry events (multiple events per ASIN). The log(FBA) column uses events with non-missing FBA on all pre- and post-window days; the log(FBM) column uses events with non-missing FBM on those days. Both entry columns require non-missing Amazon prices on all post-window days. Exit columns: for log(FBA), FBA and Amazon prices are observed each pre-period day; after exit, Amazon is absent and FBA is observed each day. For log(FBM), the analogous requirements hold for FBM.

still aim to carefully understand whether the above results can be given a causal interpretation. We hence plan to analyze trends in sales ranks and the number and types of third-party sellers before and after Amazon entry and exit events (see Fischer et al. (2025) who conduct similar analyses in the context of gasoline stations).

4.2 Asymmetric Pass-Through of Cost Shocks

We next study the pass-through margin of competition on a hybrid platform: whether Amazon and third-party sellers differ in their transmission of common cost shocks. Our empirical strategy exploits plausibly exogenous variation in cocoa-bean and coffee-bean prices, and traces these shocks into their most direct downstream retail counterparts on Amazon, namely baking cocoa and ground coffee (section 2 discusses our rationale for focusing on baking cocoa rather than other cocoa-based products). Accordingly, our pass-through estimates are reduced-form: they trace equilibrium retail price responses to plausibly exogenous commodity-cost shocks without imposing a full model of demand, seller conduct, or platform behavior. Prior work has also used coffee-bean price shocks to study pass-through; see Bettendorf and Verboven (2000), Nakamura and Zerom (2010), and Sangani (2026).

We obtain commodity price data at the daily level from Trading Economics (`tradingeconomics.com`), and (for cocoa) from the International Cocoa Organization (`www.icco.org`). See Figure 5 for plots of the corresponding prices. Reportedly, the observed hikes in price during our period of observations were due to bad weather and speculation.[15] Appendix A.1 contains further information on the respective data sources.

Figure 5: Commodity prices, cocoa beans and coffee beans.

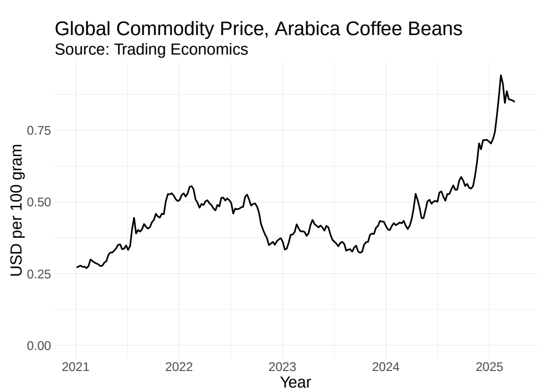

(a) Coffee.

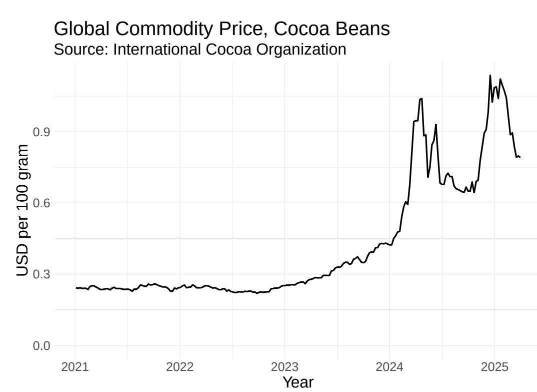

(b) Cocoa.

In order to standardize prices of products by product volume, we use OpenAI's gpt-4o-mini

[15] See `https://www.ft.com/content/4ea3116d-6a55-4d14-8a47-8988209dc1e5` and `https://www.ft.com/content/8a917bea-8ce3-11e3-ad57-00144feab7de`, both accessed 10/10/2025.

model, which we access via its API, to extract information on package sizes from product's titles. The precise procedure is described in Appendix A.2.[16]

Descriptive analysis. Panels (a) and (b) of Figure 6 present the Year-week-sellertype fixed effects estimated from regressions of prices on year-week-sellertype and manufacturer-brand fixed effects, where sellertype can be either Amazon, or third-party sellers. For cocoa, there is no statistically significant difference between Amazon's prices and those of third-party sellers in any week. For coffee, Amazon's prices are significantly lower than third-party sellers' prices, even after controlling for manufacturer-brand fixed effects. However, this discrepancy is likely driven by differences in product selection: while some third-party sellers seem to specialize in high-end brands or premium roasts, Amazon appears to focus more on standard, high-turnover brands, based on anecdotal evidence.

To better illustrate price changes of particular products, in panels (c) and (d) we plot the corresponding seller-specific chained Fisher price indices. The underlying idea of the Fisher price index, frequently used in macroeconomics, is to compare prices only for ASINs that exist in two adjacent weeks, build a weekly link (so that new or retired items do not distort the level), and chain the links across the whole period. In both categories we observe a clear upward drift in the overall price level, consistent with the large increases in input prices documented in Figure 5. The Amazon and third-party indices tend to co-move over time, indicating that, once we hold the set of products within each seller type approximately constant, third-party sellers might be raising prices somewhat more than Amazon. Since both indices are normalized to 100 in the first week of 2021, deviations between the two series at later dates can be interpreted as cumulative differences in price changes rather than in initial price levels.

Econometric specification. We estimate pass-through using a standard distributed-lag specification, which is widely used in the pass-through literature (e.g., Sangani (2026); Nakamura and Zerom (2010); Conlon and Rao (2020); Alvarez et al. (2024); amd Jean-Pierre Dubé and Gupta (2005)). This specification provides a flexible reduced-form characterization of the dynamic response of retail prices to commodity-cost shocks.

For product i observed in week t, our baseline specification is:

$$\Delta p_{it} = \sum_{k=0}^{K} \beta_k \Delta c_{t-k} + \gamma_i + \epsilon_{it} \tag{5}$$

where Δp_{it} is the change in product i's price from week t-1 to week t, measured in dollars per gram, and Δc_{t-k} is the change in the relevant commodity price, also measured in dollars per

[16] A common concern is that rising prices can affect package sizes ("shrinkflation"); however, we do not find any evidence of a change in the prevailing package size distribution over the sample period.

Figure 6: Prices over time: descriptives.

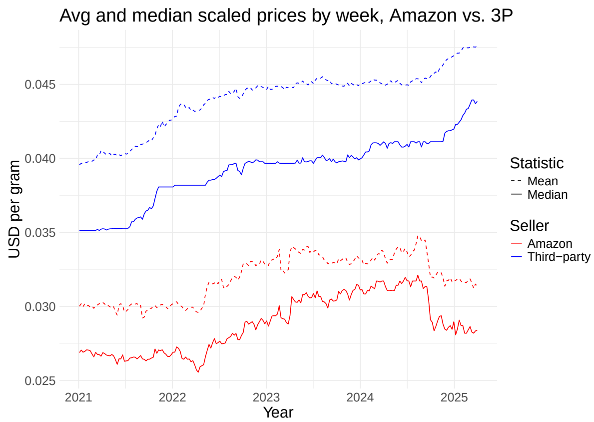

(a) Coffee, average and median prices.

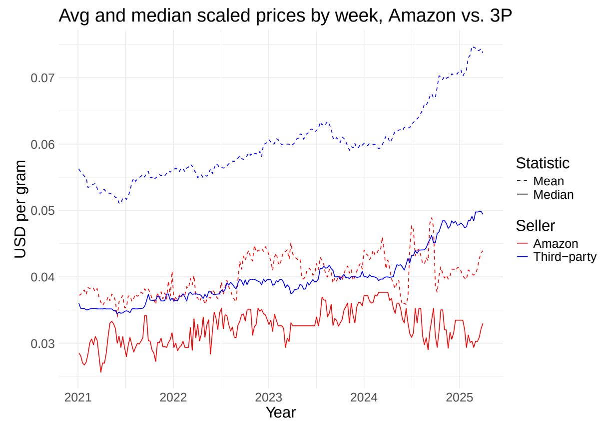

(b) Cocoa, average and median prices.

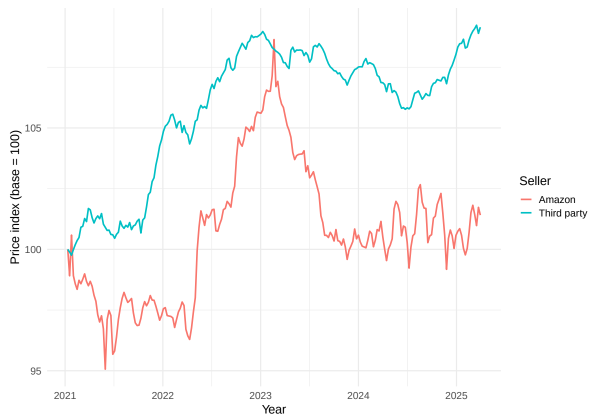

(c) Coffee, chained Fisher price index.

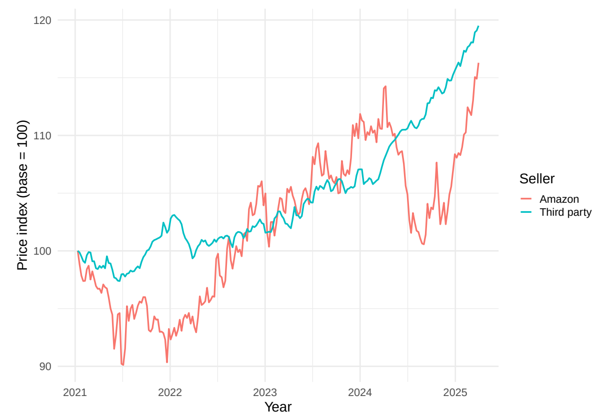

(d) Cocoa, chained Fisher price index.

Notes: Plots (a) and (b) display average and median prices for Amazon and third-party sellers, respectively. In panels (c) and (d) we plot seller–specific, chained Fisher price indices. For each seller type and pair of adjacent weeks t and $t+1$, we first restrict attention to the set of ASINs that are sold in both weeks and compute a two-period "Fisher link", defined as the geometric mean of a Laspeyres-type and a Paasche-type price relative across these overlapping products. We give equal weight to each overlapping ASIN. We then chain these weekly Fisher links over time and normalize the index to 100 in the first week of 2021, so that the resulting series traces the cumulative evolution of prices for a constant set of products within each seller type, abstracting from entry and exit in the assortment.

gram, from week $t-k$ to week $t-k-1$. Including K lags (for example, $K=6$ or $K=12$) allows us to capture delayed pass-through of upstream cost shocks. Because the dependent variable is first-differenced, time-invariant product characteristics, whether observed or unobserved, are differenced out by construction. The product fixed effects γ_i allow each ASIN to have its own average drift in price changes, which accommodates persistent differences in repricing behavior across products.

Given that both downstream prices and upstream commodity costs are measured in the same units in our data, the levels specification has a direct economic interpretation: $\sum_{k=0}^{K}\beta_k$ measures the cumulative pass-through of a change in commodity cost per gram into product prices per gram. This is also the specification we emphasize in light of recent evidence in Sangani (2026), which shows that for common commodity shocks pass-through may appear incomplete in logs even when it is close to complete in levels. We therefore use the levels specification as our baseline and report the analogous log specification in Appendix A.8.

We estimate Equation 5 using ordinary least squares. We confirm that first differences in commodity prices are stationary for both cocoa and coffee using an augmented Dickey-Fuller test. We estimate the regression separately for each product category and for prices set by Amazon itself, and for prices set by third-party sellers. We also confirmed that the distribution of package sizes of sold products does not change significantly over time, mitigating concerns of "shrinkflation" as a response to cost shocks. Appendix A.8 contains further specifications using alternative price measures.

Results. Estimates for shorter and longer-term pass-through, $\sum_{k=0}^{K} \beta_k$, of cocoa and coffee in level and with K varying between six and 48 weeks are visualized in Figure 7 and documented in Table 4. We first consider pass-through on third-party sellers' prices. We find positive and significant coefficients that are in line with prior work in economic magnitude Nakamura and Zerom (2010); Sangani (2026). Our baseline coefficients monotonically increase in magnitude as we include more lags, and approach 1 especially for cocoa when moving to a lag of 48 weeks.

We find notably different results when considering only Amazon's prices: for the case of cocoa, we obtain negative significant pass-through estimates after 24 or 48 weeks. For the case of coffee, our estimates tend to be insignificant. These findings qualitatively hold across different specifications and different sub-samples. The main noteworthy exception arises when conditioning on Buy Box prices for coffee, where we obtain mostly insignificant results in the case of third-party sellers, and a positive coefficient for a lag of 48 weeks in the case of Amazon; see Appendix A.8. This is consistent with the Buy Box algorithm being very price elastic (Lee and Musolff, 2025).

These results are in line with our theory model: While third-party sellers appear to pass through cost shocks (first partially and almost fully with longer lags), Amazon tends to either remain unresponsive to marginal cost changes or, counterintuitively, lower prices following an increase in commodity prices, at least in the time period considered.

On one hand, these results could stem from the advantage that Amazon has, as it does not need to pay itself any commission fee. However, there are admittedly other explanations for this phenomenon. For example, Amazon may seek to strengthen its reputation for offering consistently low prices (see Cabral and Xu (2021) or Cavallo, Cavallo and Rigobon (2014) for documentations of such effects). It is also thinkable that Amazon's broader portfolio and revenue

Table 4: Pass-through results: distributed lag regressions

Specification	K (lags)	Sum of Coefficients	Standard Error	T-statistic	Observations
Coffee:					
Third	6	0.18	0.06	3.11	1559229
Third	12	0.30	0.07	4.08	1559229
Third	24	0.44	0.09	4.67	1559229
Third	48	0.69	0.13	5.30	1559229
Amazon	6	-0.03	0.14	-0.20	288417
Amazon	12	-0.18	0.18	-0.97	288417
Amazon	24	-0.18	0.28	-0.63	288417
Amazon	48	0.52	0.38	1.37	288417
Cocoa:					
Third	6	0.16	0.08	1.87	149569
Third	12	0.30	0.10	3.09	149569
Third	24	0.40	0.18	2.18	149569
Third	48	0.97	0.20	4.92	149569
Amazon	6	0.10	0.24	0.42	11003
Amazon	12	0.28	0.38	0.73	11003
Amazon	24	-1.32	0.46	-2.88	11003
Amazon	48	-1.57	0.72	-2.19	11003

Notes: Each row reports the cumulative pass-through estimate, $\sum_{k=0}^{K} \beta_k$, from the distributed-lag specification in equation 5, where K denotes the number of weekly lags of the commodity-price change (cocoa or coffee, respectively) included in the regression. Regressions are estimated separately for third-party seller prices (*Third*) and Amazon prices (*Amazon*) using the baseline levels specification. The dependent variable is the weekly change in product price per gram, and the regressors are the contemporaneous and lagged weekly changes in the commodity price per gram. The sample is a weekly ASIN panel of products from January 2021 to March 2025. Standard errors are two-way cluster-robust. T-statistics are computed using these standard errors.

Figure 7: Pass-through estimates from distributed-lag regressions (Table 4)

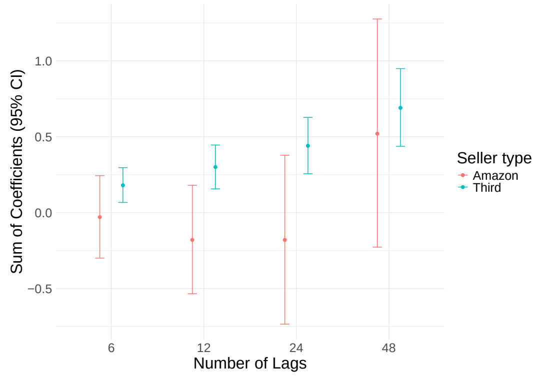

(a) Coffee.

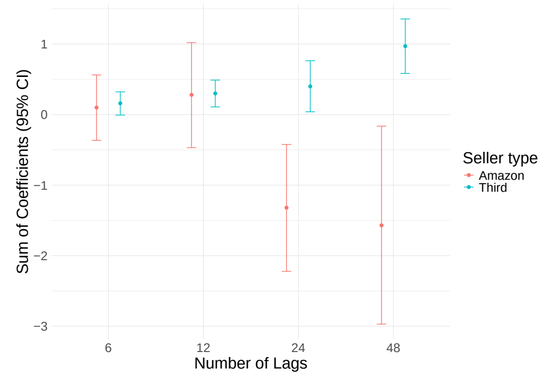

(b) Cocoa.

streams can enable it to cross-subsidize specific items, selling them below costs or very small margins, or that it is more efficient.

5 Future Work: Generalized Model

While our theoretical model isolates the core incentive mechanisms, i.e., the tension between retail revenue and commission fee income, its single-product assumption limits our ability to perform counterfactual policy analysis. To bridge the gap between theory and the complex reality of the marketplace, we propose extending our framework to a generalized quantitative model. This extension rests on two structural pillars:

First, we relax the single-product restriction by allowing each seller $s \in S$ (including the platform) to offer a portfolio of goods $g \in G$. The relevant demand system is thus defined over the space of *varieties*, where a variety is a unique seller-product pair. The total number of potential varieties is $N_S \times N_G$. Empirically, we leverage the fact that this matrix is highly sparse—most sellers list only a limited subset of available goods—to ensure computational feasibility. By defining profit maximization over the portfolio of varieties, this setup allows us to capture cross-product pricing externalities. This is particularly crucial for the platform, which may price a subset of products aggressively (as suggested by our Fact 3) to drive traffic to the platform, thereby internalizing cross-category spillovers, a mechanism our current model abstracts from.

Second, to mirror the consumer's actual decision process on Amazon, we model demand using a two-level nested logit framework. Consumers first select a specific product (ASIN) based on product differentiation and an inclusive value index that aggregates the attractiveness of all offers for that product (Stage 1: Shelf Competition). Then, conditional on the ASIN choice, consumers select a specific seller (Stage 2: Buy Box Competition). This structure implies a realistic substitution pattern: competition is intense within an ASIN (where sellers are close substitutes) but dampened across ASINs. In addition, the lower nest (Buy Box competition) allows for a structural estimation of self-preferencing. By parameterizing the seller-specific utility $\delta_k = \xi_j + \beta \cdot \mathbb{I}(\text{Amazon}_k) - \alpha p_k$, we can model the "Amazon bias" (β) as the extent to which the platform steers demand towards its own offers controlling for price and observable quality.

This structural framework enables us to decompose consumer welfare changes into variety effects, price effects, and efficiency gains, providing a rigorous basis for quantitative evaluation and counterfactual analysis.

6 Conclusion

This paper identifies a distinct pricing mechanism on dual-mode e-commerce platforms: because the platform both earns commissions on third-party sales and retails on its own marketplace, its pricing incentives differ from those of third-party sellers. In a stylized model, this generates two key implications: platform retail presence can either soften or intensify price competition, and common cost shocks need not be passed through symmetrically across seller types. Using data from Amazon U.S., we document reduced-form evidence consistent with both implications. Third-party prices fall following Amazon entry and rise following Amazon exit, and third-party sellers appear to pass through commodity cost shocks more strongly than Amazon. These findings highlight a novel channel through which dual-mode operation may shape competition on e-commerce platforms.

Our results are relevant not only for our understanding on price competition on dual-mode platforms, but also in light of recent antitrust cases that focus on pricing practices and restraints on third-party prices in Germany and the U.S. (Federal Trade Commission, 2024; Bundeskartellamt, 2025).

References

Alvarez, Santiago E., Alberto Cavallo, Alexander MacKay, and Paolo Mengano (2024) "Markups and Cost Pass-through Along the Supply Chain," *Working Paper*, https://alexandermackay.org/files/Markups%20and%20Cost%20Pass-through%20Along%20the%20Supply%20Chain.pdf.

Anderson, Simon and Özlem Bedre Defolie (2024) "Hybrid platform model: monopolistic competition and a dominant firm," *The RAND Journal of Economics*, 55 (4), 684–718, https://doi.org/10.1111/1756-2171.12478.

Arcidiacono, Peter, Paul B. Ellickson, Carl F. Mela, and John D. Singleton (2020) "The Competitive Effects of Entry: Evidence from Supercenter Expansion," *American Economic Journal: Applied Economics*, 12 (3), 175–206, 10.1257/app.20180047.

Bennati, Luca (2026) "Vertical Integration and E-commerce Competition: Evidence from Amazon Marketplace," *Working Paper*, https://drive.google.com/file/d/1GCzljX1haX5pLL_IXGlQ4r7fWgCJI3ee/view.

Bettendorf, L and F Verboven (2000) "Incomplete transmission of coffee bean prices: evidence from The Netherlands," *European Review of Agricultural Economics*, 27 (1), 1–16, 10.1093/erae/27.1.1.

Bisceglia, Michele and Jorge Padilla (2023) "On sellers' cooperation in hybrid marketplaces," *Journal of Economics & Management Strategy*, 32 (1), 207–222.

Boldrini, Michela and Francesco Clavorà Braulin (2025) "Taming Tech Giants' Algorithms: An analysis of consumers' choices on Amazon,"Technical report, https://drive.google.com/file/d/1nmJjwZF3mRezOn2T-1zGUNeBQg4OGAQ1/view, Mimeo.

Breslin, Seamus (2025) "How to Win the Amazon Buy Box," February, https://www.repricerexpress.com/win-amazon-buy-box/, Accessed: 2026-02-12.

Bundeskartellamt (2025) "Bundeskartellamt has concerns about Amazon's use of so-called price control mechanisms," https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2025/2025_06_02_Amazon.html, June, Press release.

Cabral, Luís and Lei Xu (2021) "Seller reputation and price gouging: Evidence from the COVID-19 pandemic," *Economic Inquiry*, 59 (3), 867–879, https://doi.org/10.1111/ecin.12993.

Cavallo, Alberto (2019) "More Amazon effects: online competition and pricing behaviors," *Jackson Hole Economic Symposium Conference Proceedings*.

Cavallo, Alberto, Eduardo Cavallo, and Roberto Rigobon (2014) "Prices and Supply Disruptions during Natural Disasters," *Review of Income and Wealth*, 60 (S2), S449–S471, https://doi.org/10.1111/roiw.12141.

Chen, Le, Alan Mislove, and Christo Wilson (2016) "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace," *Proceedings of the 25th international conference on World Wide Web*, 1339–1349.

Chen, Nan and Hsin-Tien Tsai (2024) "Steering via algorithmic recommendations," *The RAND Journal of Economics*, 55 (4), 501–518, https://doi.org/10.1111/1756-2171.12481.

Conlon, Christopher T. and Nirupama L. Rao (2020) "Discrete Prices and the Incidence and Efficiency of Excise Taxes," *American Economic Journal: Economic Policy*, 12 (4), 111–143, 10.1257/pol.20180658.

de Cornière, Alexandre and Greg Taylor (2014) "Integration and search engine bias," *The RAND Journal of Economics*, 45, 576–597, 10.1111/1756-2171.12063.

Crawford, Gregory S, Matteo Courthoud, Regina Seibel, and Simon Zuzek (2022) "Amazon entry on Amazon marketplace," *CEPR Discussion Paper No. 17531*.

Cure, Morgan, Matthias Hunold, Reinhold Kesler, Ulrich Laitenberger, and Thomas Larrieu (2022) "Vertical integration of platforms and product prominence," *Quantitative Marketing and Economics*, 20, 353–395.

Dendorfer, Florian and Regina Seibel (2026) "What's In the Box? The Effect Of Self-Preferencing On Amazon Marketplace,"Technical report, `https://reginaseibel.github.io/publication/amazon_box/amazon_box.pdf`, Mimeo.

Drugov, Mikhail and Doh-Shin Jeon (2019) "Vertical integration and algorithm bias," Mimeo.

Etro, Federico (2021) "Product selection in online marketplaces," *Journal of Economics & Management Strategy*, 30, 614–637, 10.1111/jems.12428.

——— (2023) "Hybrid Marketplaces with Free Entry of Sellers," *Review of Industrial Organization*, 62, 119–148.

Fabra, Natalia and Mar Reguant (2014) "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, 104 (9), 2872–99.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay (2023) "Self-Preferencing at Amazon: Evidence from Search Rankings," *AEA Papers and Proceedings*, 113, 239–43, 10.1257/pandp.20231068.

——— (2025) "Vertical Integration and Consumer Choice: Evidence from a Field Experiment," NBER Working Paper 34135, National Bureau of Economic Research, 10.3386/w34135.

Federal Trade Commission (2024) "Amended Complaint: Amazon.com, Inc. (Amazon eCommerce)," https://www.ftc.gov/system/files/ftc_gov/pdf/01712024.03.14RedactedAmendedComplaint%20%28002%29.pdf.

Fischer, Kai, Simon Martin, and Philipp Schmidt-Dengler (2025) "The Heterogeneous Effects of Entry on Prices,"Technical Report 10566, https://drive.google.com/file/d/1kMBYUHED-_sU9xMH76LaQerI-Mi5V_Uc/view, Forthcoming in Journal of the European Economic Association.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker (2020) "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, 12 (2), 303–342, https://www.jstor.org/stable/26909455.

Genakos, Christos and Mario Pagliero (2022) "Competition and Pass-Through: Evidence from Isolated Markets," *American Economic Journal: Applied Economics*, 14 (4), 35–57.

Goolsbee, Austan D. and Peter J. Klenow (2018) "Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce," *AEA Papers and Proceedings*, 108, 488–92.

Gorodnichenko, Yuriy and Oleksandr Talavera (2017) "Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration," *American Economic Review*, 107 (1), 249–282.

Gutiérrez, Germán (2022) "The Welfare Consequences of Regulating Amazon," http://germangutierrezg.com/Gutierrez2021_AMZ_welfare.pdf.

Hagiu, Andrei, Tat-How Teh, and Julian Wright (2022) "Should platforms be allowed to sell on their own marketplaces?" *The RAND Journal of Economics*, 53 (2), 297–327.

He, Leshui, Imke Reimers, and Benjamin Shiller (2022) "Does Amazon Exercise Its Market Power? Evidence from Toys "R" Us," *The Journal of Law and Economics*, 65 (4), 665–685.

He, Sherry, Brett Hollenbeck, and Davide Proserpio (2022) "The Market for Fake Reviews," *Marketing Science*, 41 (5), 896–921.

amd Jean-Pierre Dubé, David Besanko and Sachin Gupta (2005) "Own-Brand and Cross-Brand Retail Pass-Through," *Marketing Science*, 24 (1), 123–137.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard (2019) "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 11 (4), 181–209, 10.1257/app.20180010.

Lee, Kwok Hao and Leon Musolff (2025) "Two-Sided Markets Shaped by Platform-Guided Search," *Working Paper*, https://lmusolff.com/papers/TwoSidedMarketsPlatformGuidedSearch.pdf.

Madsen, Erik and Nikhil Vellodi (2023) "Insider Imitation," *Working Paper*.

Miller, Nathan H., Matthew Osborne, and Gloria Sheu (2017) "Pass-through in a Concentrated Industry: Empirical Evidence and Regulatory Implications," *The RAND Journal of Economics*, 48 (1), 69–93, http://www.jstor.org/stable/26305442.

Montag, Felix, Robin Mamrak, Alina Sagimuldina, and Monika Schnitzer (2023) "Imperfect Price Information, Market Power, and Tax Pass-Through," *George J. Stigler Center for the Study of the Economy & the State* (Working Paper # 337).

Nakamura, Emi and Dawit Zerom (2010) "Accounting for Incomplete Pass-Through," *The Review of Economic Studies*, 77 (3), 1192–1230, https://doi.org/10.1111/j.1467-937X.2009.589.x.

Pless, Jacquelyn and Arthur A. van Benthem (2019) "Pass-Through as a Test for Market Power: An Application to Solar Subsidies," *American Economic Journal: Applied Economics*, 11 (4), 367–401.

Sangani, Kunal (2026) "Pass-Through in Levels and the Unequal Incidence of Commodity Shocks," *forthcoming in the Quarterly Journal of Economics*, https://kunalsangani.com/files/complete_passthrough_live.pdf.

Weyl, E. Glen and Michal Fabinger (2013) "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 121 (3), 528–583.

Zhu, Feng and Qihong Liu (2018) "Competing with complementors: An empirical look at Amazon.com," *Strategic Management Journal*, 39 (10), 2618–2642, https://doi.org/10.1002/smj.2932.

A Empirical Appendix

A.1 Additional information on commodity data

Cocoa beans. The cocoa bean pricing data we use stems from the International Cocoa Organization (ICCO) and is available via www.icco.org/statistics. We use the ICCO daily price for cocoa beans in US$, defined as "the average of the quotations of the nearest three active futures trading months on ICE Futures Europe (London) and ICE Futures US (New York) at the time of London close". We refer the interested reader to the website for further information.

Arabica coffee beans. Daily pricing data on Arabica coffee beans are obtained from Trading Economics, available on tradingeconomics.com/commodity/coffee. The data corresponds to Arabica Coffee C Futures contracts that trade on the Intercontinental Exchange (see https://www.ice.com/products/15/Coffee-C-Futures), and is the global benchmark for Arabica Coffee.

For both price series, we create weekly average prices by averaging across all observed prices in a given week. Figure 5 shows the weekly time series of both prices.

A.2 Data Collection from Keepa and Creation of Daily Product Panel

As a first step, we use Keepa's "Product Finder" tool to arrive at a list of ASINs related to a given product category (e.g., Coffee or Men's T-Shirts). Table 5 shows the query parameters used for each of the products that were queried. Note that a product (ASIN) on Amazon can belong to one *Root Category*, and to multiple *Browse Nodes*[17]. The latter are called *Subcategories* on the Keepa website. We configure the Keepa Product Finder filters so that only subcategories directly related to our target product (e.g., "Ground Coffee") are included, while certain adjacent but irrelevant subcategories (e.g., "Instant Coffee") may explicitly be excluded.[18]

For product categories with over 10,000 products listed on amazon.us, we only collect data on those 10,000 products with the lowest current Sales Rank or the lowest 180-day average sales rank at the time of scraping, thereby focusing our attention on products that actually generate sales and are (normally) in stock. After obtaining a list of ASINs for a given product category, we query these ASINs one-by-one through by doing "Product Object" queries[19] through Keepa's Application Programming Interface (API), and thereby obtain product and price information. We query all categories at multiple instances in time in the period of July 2024 through March 2026.

[17]See https://webservices.amazon.com/paapi5/documentation/use-cases/organization-of-items-on-amazon/browse-nodes/browse-nodes-and-items.html (accessed on 25/11/2024). The Browse Node IDs of the subcategories listed in Table 5 are not displayed for readability reasons, but can be provided upon request.

[18]This way, we can focus on a set of truly highly substitutable products and exclude products that belong to different, dis-similar categories (possibly as a result of misclassification).

[19]See https://keepa.com/#!discuss/t/product-object/116 (last accesed 18/02/2026).

Table 5: Query parameters used in initial product data collection process on Keepa.

Root Category and Product Category	Included and excluded subcategories
Root Category: Grocery & Gourmet Food	
Coffee	· include: Ground Coffee · exclude: Tea; Instant Coffee
Cocoa	· include: Cocoa · exclude: Chocolate Drink Mixes; Hot Chocolate & Malted Drinks; Thanksgiving, Tea & Cocoa
Pasta	· include: Pasta & Noodles · exclude: Fresh Pasta & Sauces; Gnocchi
Root Category: Electronics	
TVs	· include: LED & LCD TVs · exclude: Remote Controls; TV Mounts, Stands & Turntables; Satellite Equipment
USB Hubs	· include: USB Hubs
Micro SD Cards	· include: Micro SD Cards
Root Category: Appliances	
Washing Machines	· include: Washers · exclude: Countertop Dishwashers; Built-in Dishwashers; Portable Dishwashers; Washer Parts & Accessories; Portable Washers
Root Category: Clothing, Shoes & Jewelry	
Men's T-Shirts	· include: Men's T-Shirts · exclude: Tights, Pants & Shorts; Leggings

After collecting the list of ASINs, we exclude products which (either based on manual inspection or based on the variable "type") seem to belong to a different category and are thus misclassified. For products in the Pasta category, we exclude products marked with terms like "buckwheat", "rice" or "gluten-free", so as to focus on products made from wheat; for products in the coffee category, we exclude chicory or instant coffee sticks, for instance. Within a given category, we only use data on products that are ever on sale on amazon.com and thus have a price indicated at any point since 2021.

The original "Product Object" data downloaded through the Keepa API essentially consists of separate datasets for the lowest "marketplace" offer (which can be Amazon or a third-party seller's price, called NEW in Keepa); the lowest price set by any third-party seller price operating under FBA; the lowest price set by any third-party seller price operating under FBM; Amazon's price; and the price of the seller whose offer is displayed in the Buybox. One row in this original data corresponds to a given price change at a given moment in time (a price becoming unavailable or available corresponds to a "change" as well). We define the lowest third-party price (NEW_third) as the marketplace price (NEW) whenever it differs from the Amazon price, and as the minimum of the lowest FBA and FBM prices otherwise.

Adjustment for occasionally occurring delayed Amazon price recording. The product-level price data from Keepa occasionally exhibits a timing discrepancy in which Amazon's offer is detected with a short lag. Specifically, when a previously out-of-stock Amazon offer returns to the marketplace, Keepa's system may register the resulting change in the marketplace price (NEW) and seller count (COUNT_NEW) before identifying the Amazon offer itself. During this interim period – lasting only a few days – the AMAZON price field remains missing even though the marketplace price already reflects Amazon's presence, which would cause the lowest third-party price (NEW_third) to be incorrectly attributed. We confirmed this data-collection artifact directly with the Keepa team[20]. To address it, we apply a conservative correction procedure that detects instances in which (i) the seller count increases by exactly one, (ii) the marketplace price simultaneously decreases, (iii) the Amazon price is not yet recorded, (iv) no FBA or FBM price equals the marketplace price on the affected days (ruling out the possibility that the price drop is attributable to a third-party seller), and (v) within seven days the Amazon price appears while the seller count remains *unchanged*. When all four conditions are met, we backfill the Amazon price to the date of the initial seller-count increase. This adjustment affects a small number of product-days (0.001% in the case of Cocoa and 0.002% in the case of coffee, for instance) and ensures that Amazon's entry date is correctly recorded.

Dealing with a small number of apparent "pricing errors". We noticed a certain number of apparent "pricing errors" in the original data, i.e., instances where a product's price jumps up by several multiples of the original price, before coming down. These can be due to, e.g., data entry mistakes, bot malfunctions. For instance, a package of coffee bearing the ASIN B09B8TP9RQ, that is typically priced at \$24.99, is suddenly offered at \$5,889.34 on February 20th to 26th, 2025 – an almost 200-fold increase in its price – before coming back down. We never observe such patterns in the case of Amazon's offers, only for third-party sellers' prices. To detect these upward price deviations in a conservative, but sensible way and exclude them from our dataset upfront, we employ a multi-stage procedure to identify these errors at both the observation and product levels. For each product (ASIN), we calculate the interquartile range (IQR = Q75 - Q25) of observed prices. An individual price observation is flagged as an outlier if it satisfies all three conditions: (1) Price $>$ Q75 + 30 $\times$ IQR; (2) Price - Q75 $>$ 2 $\times$ median_all, where median_all is the median price across all products; and (3) Price $>$ 2 $\times$ median_product. These criteria jointly ensure that flagged observations represent extreme deviations in both product-specific terms (IQR criterion), absolute terms (preventing spurious flags in low-variance products), and relative terms (substantial deviation from the product's typical price level). Next, we allow for the exclusion of entire *products* from the sample when outlier prevalence suggests systematic data

[20] See https://keepa.com/#!discuss/t/price-update-frequency-for-amazon/18645 (last accessed 18/02/2026).

quality issues.[21] Note that, in the case of coffee or cocoa, there are no such systematic data quality issues for any products in our sample. More generally, only around 0.2% of all ASINs are affected by this procedure. Finally, in the coffee category, we exclude one brand ("San Marco Coffee") which operated under a dropshipping arrangement until 2024. For this brand, the data show an extreme repricing event, followed by a synchronized exit of Amazon offering any products by that brand. The observed patterns are consistent with a termination in the vendor relationship, which was confirmed to us by the brand, and thus is likely independent from any pass-through of cost shocks.

After creating a daily panel of products' prices (recorded at midnight of a given date), we further exclude products if their average third-party sellers' (Amazon's) price (standardized by weight when considering grocery products) across all observed days lies above the 99th, or below the 1st percentile of all third-party sellers' (Amazon's) daily prices: these products are likely non-coffee products that are misclassified, and should hence not be part of our sample.

A.3 Extracting Relevant Product Information Using OpenAI API

In our analysis, precise extraction of product package sizes is crucial. However, package sizes from the relevant variables included by Keepa are missing for many of the products. Instead, package sizes are being accurately described in the product title, but due to nonstandardized formatting, are difficult to extract using standard string-matching methods. To address this, we employ OpenAI's gpt-4o-mini model. The prompt instructs the model to identify and compute information on total weight or volume of the content as indicated in the product title, and is displayed below:

```
You are tasked with extracting the total weight or volume of a grocery
   product from its description. Follow these specific rules:

    1. **Accurate extraction only**: Use only the information explicitly
       provided in the product description. Do not infer or assume any
       details.
    2. **Output format**: Provide the result as ‘number|unit‘.
       - If the description includes a total weight or volume (including
          multipacks), calculate the combined total if necessary.
       - Units like ounces (oz), grams (g), pounds (lb), kilograms (kg),
          liters (L), or milliliters (ml) should be preserved as stated.
    3. **No weight or volume available**: If no weight or volume
       information is provided, return ‘NA|NA‘.
```

[21] Specifically, we remove all observations for a product if either: (1) more than 30% of observed prices are flagged as outliers, or (2) the product has at least 10 outliers and fewer than 50 total price observations.

```
**Examples**:

- **Input**: "Starbucks Caffe Verona Coffee, Dark, Ground, 12-Ounce
   Bags (Pack of 3)"
  **Output**: 36|oz

- **Input**: "Four Sigmatic Think Mushroom Coffee | Organic Ground
   Coffee with Lion's Mane Mushroom and Chaga Mushroom | 12oz Bag"
  **Output**: 12|oz

- **Input**: "AmazonFresh Go For The Bold Ground Coffee, Dark Roast"
  **Output**: NA|NA

- **Input**: "AmazonFresh French Vanilla Flavored Coffee, Ground,
   Medium Roast, 500g"
  **Output**: 500|g

Now, extract the total weight or volume from the following description
   :

---
{text}
---
```

To guarantee consistency and reproducibility, we set the temperature parameter to 0. We moreover include a system message – "role": "system", "content": "You are an expert at analyzing grocery product descriptions and extracting accurate weight or volume information. You strictly follow instructions, avoid assumptions, and format answers as specified." – which primes the model as a domain expert and ensures it adheres rigorously to the extraction protocol without introducing unwarranted assumptions.

Table 6 displays the resulting distribution of weights across ASINs.

Table 6: Distribution of product weights across ASINs in grocery categories

Category	ASINs	ASINs with weight	Mean	Median	SD	P10	P90
Cocoa	1849	1802	1258.3	453.6	3147.7	150.0	2268.0
Coffee	17250	16981	904.8	453.6	1426.6	250.0	2041.2
Pasta	8908	8673	2063.4	997.9	7564.9	340.0	5443.1

Note: The table reports ASIN-level weight statistics (in grams) for Cocoa, Coffee, and Pasta products.

A.4 Stylized Facts of Dual-Mode Selling: Further Evidence

Figure 8: Extent of dual-mode selling across categories, popular products

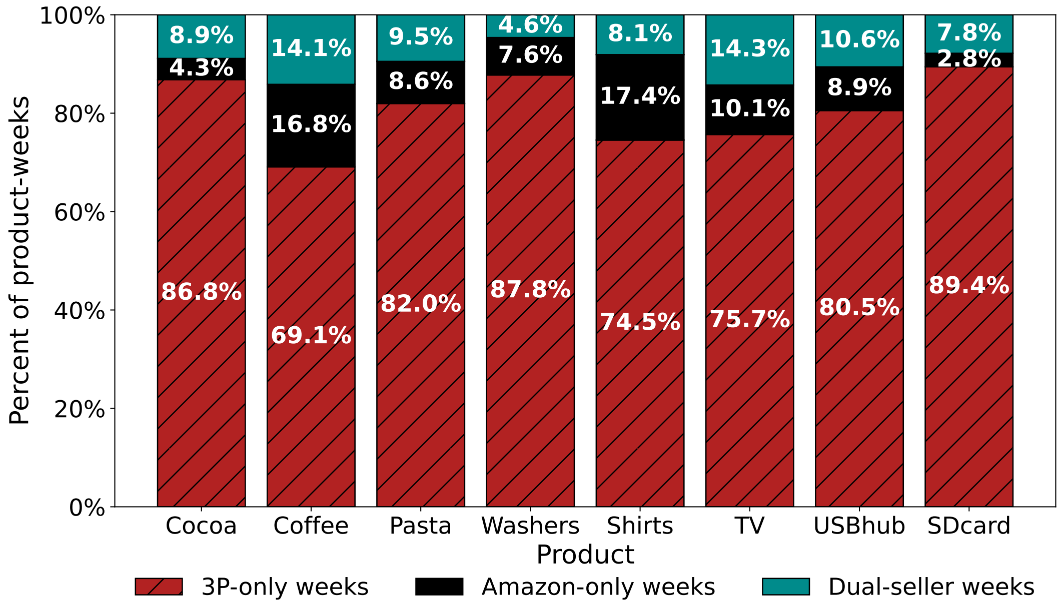

Notes: Bars show, for each category, the share of ASIN-week observations in which the product is sold only by third-party sellers, only by Amazon, or by both. The sample contains **only popular** ASINs (as defined in Section 2.2, observed weekly from January 2021 through December 2024. A product-week is classified as 3P-only (Amazon-only) if at least one 3P (Amazon) offer is recorded and the other side is absent; it is classified as dual-seller if both have offers in the same week.

A.5 Analysis of Prices: Further Evidence

Price adjustment frequency. Note that when broadening the sample to include all active listings regardless of which seller types are present in a given week, a divergence in repricing frequency emerges. As shown in Figure 10, the share of products that record no price change from one week to the next as a share of all products with available prices in both weeks exceeds 70% for third-party sellers, whereas for Amazon this share varies between 30% and 65% across both cocoa and coffee categories. This pattern holds for the subset of popular products as well. Since Fact 4 establishes that no such difference exists on jointly-offered products, this aggregate divergence must be driven by products sold exclusively by one seller type — that is, Amazon-exclusive products adjust prices more frequently than 3P-exclusive products, on average.

Table 7: Price differences for identical products: Coffee

	All products			Popular products		
Dependent Variable:	log(price)					
Compared to	Any 3P	FBA	FBM	Any 3P	FBA	FBM
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
amazon	-0.266***	-0.191***	-0.424***	-0.262***	-0.198***	-0.426***
	(0.007)	(0.008)	(0.007)	(0.008)	(0.009)	(0.007)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearweek	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	303,704	214,578	289,984	262,784	187,936	251,722
Adjusted R^2	0.88558	0.92299	0.91656	0.88583	0.92242	0.91789

Notes: The unit of observation is an ASIN–week in the Coffee category. The dependent variable is the log of the posted price. For each column, we restrict the sample to ASIN–week observations in which both Amazon and at least one third-party seller of the relevant type (any 3P, FBA, or FBM) are active. The coefficient on amazon measures the average log price difference between Amazon and the corresponding third-party offer for the same product in the same week. "Popular products" are defined above. Standard errors, clustered at the ASIN level, are reported in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Table 8: Price differences for identical products: Cocoa

	All products			Popular products		
Compared to	Any 3P	FBA	FBM	Any 3P	FBA	FBM
Dependent Variable:	log(price)					
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
amazon	-0.038	-0.058*	-0.265***	-0.032	-0.046	-0.277***
	(0.046)	(0.030)	(0.035)	(0.052)	(0.032)	(0.038)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearweek	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	13,672	10,469	13,081	11,886	9,282	11,387
Adjusted R^2	0.90262	0.93962	0.93003	0.87095	0.92027	0.90622

Notes: The unit of observation is an ASIN–week in the Cocoa category. The dependent variable is the log of the posted price. For each column, we restrict the sample to ASIN–week observations in which both Amazon and at least one third-party seller of the relevant type (any 3P, FBA, or FBM) are active. The coefficient on amazon measures the average log price difference between Amazon and the corresponding third-party offer for the same product in the same week. "Popular products" are defined above. Standard errors, clustered at the ASIN level, are reported in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Table 9: Price differences for identical products: Men's T-shirts

	All products			Popular products		
Compared to	Any 3P	FBA	FBM	Any 3P	FBA	FBM
Dependent Variable:			ln_price			
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
amazon	-0.002	-1.7×10^{-5}	-0.123***	-0.003	0.0005	-0.206***
	(0.002)	(0.0009)	(0.013)	(0.002)	(0.0009)	(0.021)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearweek	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	294,256	281,086	168,745	222,880	218,194	120,576
Adjusted R^2	0.85209	0.89585	0.85121	0.86976	0.90314	0.88597

Clustered (asin) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 10: Price differences for identical products: SD Cards

	All products			Popular products		
Compared to	Any 3P	FBA	FBM	Any 3P	FBA	FBM
Dependent Variable:			ln_price			
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
amazon	-0.016	0.019	-0.095***	0.011	0.091**	-0.079**
	(0.021)	(0.025)	(0.023)	(0.032)	(0.037)	(0.036)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearweek	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	22,702	17,569	21,465	12,802	9,546	12,068
Adjusted R^2	0.91841	0.93266	0.91566	0.90487	0.92108	0.90090

Clustered (asin) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Figure 9: Competition by seller regime

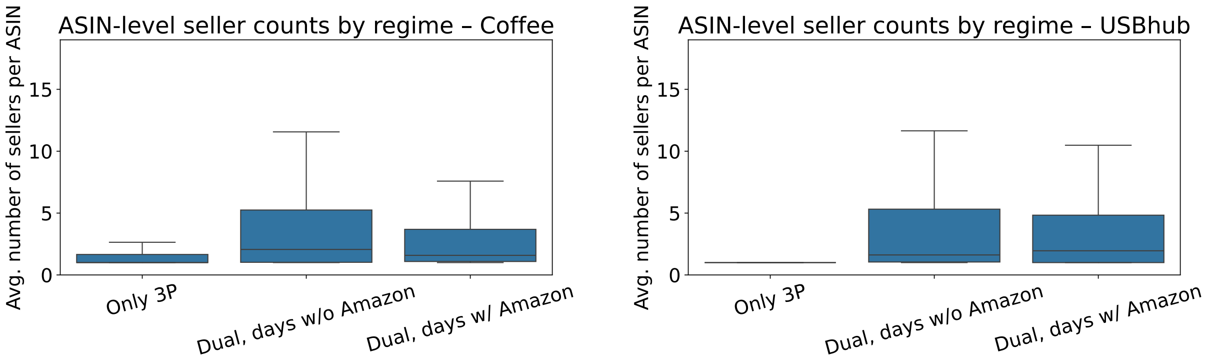

(a) Coffee. (b) USB Hubs.

Notes: The figure plots, for the Coffee and USB hub categories, the distribution across ASINs of the average daily number of sellers in three regimes. "Only 3P" includes ASINs that are never supplied by Amazon but have at least one 3P seller; for each such ASIN we average the number of sellers over days on which a 3P seller is present. "Dual, days w/o Amazon" restricts to ASINs that are supplied at some point by both Amazon and 3P sellers, but averages the number of sellers only over days when Amazon is absent and at least one 3P seller is present. "Dual, days w/ Amazon" uses the same set of dual ASINs but averages over all days when Amazon is present, regardless of whether a 3P seller is active. (We restrict to days with at least one seller present.)

Figure 10: Percentage of products that recorded no price change from week $t-1$ to week t

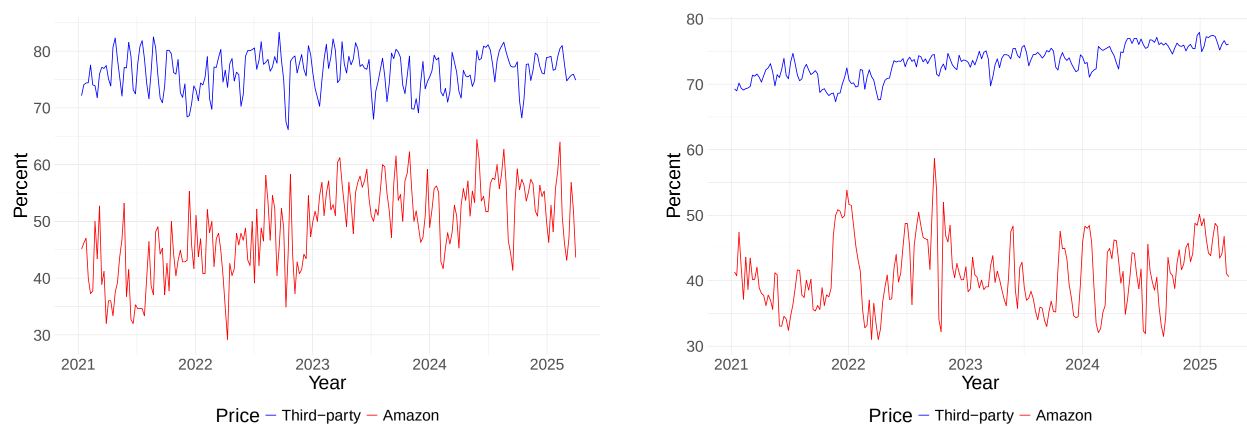

(a) Cocoa. (b) Coffee.

Notes: Both plots show the percentage of products that recorded no price change from week $t-1$ to week t out of all products that have prices available in week $t-1$ and week t.

A.6 What Determines Amazon's Decision to Sell a Product as a Retailer?

As Figure 2 shows, the presence of Amazon as a seller (in terms of number of products) varies across product categories. What may explain Amazon's decision to begin entering a given ASIN as a re-seller?

Table 11: Price differences for identical products: USB hubs

	All products			Popular products		
Compared to	Any 3P	FBA	FBM	Any 3P	FBA	FBM
Dependent Variable:	ln_price					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
amazon	0.003	0.028	-0.090***	0.005	0.053	-0.094***
	(0.017)	(0.021)	(0.014)	(0.021)	(0.033)	(0.018)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearweek	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	37,110	27,557	32,355	20,194	15,013	17,299
Adjusted R^2	0.92869	0.93943	0.94295	0.93098	0.93554	0.94917

Clustered (asin) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

A.6.1 Survey of Previous Literature

Several academic papers shed light on this question by studying the predictors of Amazon's decision to directly re-sell a product (i.e., to "enter").

Zhu and Liu (2018) employs data from 2013 and 2014 and finds that Amazon is more likely to begin re-selling products that experience higher sales and have better reviews, and which are not yet using Amazon's fulfillment services.

A working paper by Crawford et al. (2022) uses proprietary data on prices and revenues in Amazon's Home & Kitchen category. The paper focuses on Amazon's decision to begin selling products for which only a single third-party merchant supplied that product in the first four weeks of the product's existence. The authors find that Amazon starts re-selling products that experience high demand growth and where competition is low (i.e., few active sellers are present). This is true both relative to products which Amazon does *not* enter, as well as relative to the entry decisions by big (high-revenue) third-party sellers. Entry by Amazon occurs even if a large or efficient third-party seller is present. Amazon is also more likely to enter when third-party listings exist but generate no sales, suggesting the presence of low-quality or capacity-constrained sellers. The authors argue that these patterns suggest the internalization of platform externalities as a possible motive for Amazon's entry decisions. Among those externalities are expanding product variety, ensuring product availability, and mitigating seller market power.

Concerning Amazon's presence as a re-seller more generally, the authors find that nearly half of the revenue from products in which Amazon is a reseller comes from cases of *de novo* entry – products that had not been sold by any third-party seller before Amazon's entry as a reseller.

Amazon's private label products play a minor role in the Kitchen & Home category, accounting for only 0.3% of total revenues. Amazon is present in products accounting for 39% of revenues in the Home & Kitchen category.

The findings by Bennati (2026), who studies the market for headphones and uses sales estimates stemming from the market intelligence company AmzScout, are very much in line with this. The author's findings suggest that third-party sellers concentrate on products in the long tail, whereas Amazon re-sells products that are in higher demand and of higher quality.

A.6.2 Amazon Entry in Our Data

Table 12: Correlates of whether a product is ever supplied by Amazon, Coffee products only.

	Dependent variable:	
	$\mathbb{1}$\{product ever sold by Amazon\}	
1\{popular\}	0.223**	0.162***
	(0.090)	(0.047)
1\{has any reviews\}	0.175***	0.025
	(0.061)	(0.033)
avg number of 3P sellers	0.008	0.007***
	(0.013)	(0.003)
package size (weight, in kg)	$1.002\ e^{-8}$	$8.316\ e^{-10}$
	$(1.913\ e^{-8})$	(5.273ee−9)
product was ever stocked out	0.164***	0.014
	(0.056)	(0.031)
Manufacturer FE		✓
Brand FE		✓
Observations	12,759	12,664
Adjusted R^2	0.076	0.788

SEs clustered at manufacturer level. *p<0.1; **p<0.05; ***p<0.01

Notes: Cross-sectional dataset was created by beginning with a daily panel ranging from January 1st 2021 to March 31 2025, and collapsing it to the cross-sectional level. 1\{popular\} is the popularity measure defined in Section 2.2, indicating the proportion of days on which a product's sales rank is below the median sales rank across products, and hence varies between 0 and 1. The average number of 3P sellers is defined as an average across all days on which a product is available for sale. The coefficient of a linear regression are shown, and each observation is a product.

Taking the example of coffee, we find that, in the cross-section, the brand and the manufacturer are main factors correlated with Amazon's entry patterns. Controlling for brand and manufacturer, we moreover find that products where Amazon is a reseller at any point during our period of observation are likely to derive more sales, and are likely to feature more other sellers, as can be seen in Table 12. The finding that more competition by third-party sellers is correlated with Amazon re-selling the product could in fact be driven by unobserved (by the econometrician) demand shocks, which might cause entry by both Amazon and other merchants. Hence, our measure of "high competition" might in fact be a proxy for "high demand". (Of course, this result can by no means be interpreted causally, and is purely descriptive.)

In the time series, interestingly we find that, across our product categories, the overall number of products re-sold by Amazon has been constant over the years, while the number of products re-sold by third-party sellers has increased over time (not shown in paper).

A.7 Amazon Retail Presence and Third-Party Price Levels: Background and Additional Specifications

Section A.7.1 provides descriptives on the sample used, while Section A.7.2 shows additional specifications.

A.7.1 Background

We here show descriptive facts on the sample used to better understand Amazon's varying retail presence.

Figure 11: Qualifying event windows (entry and exit) over time

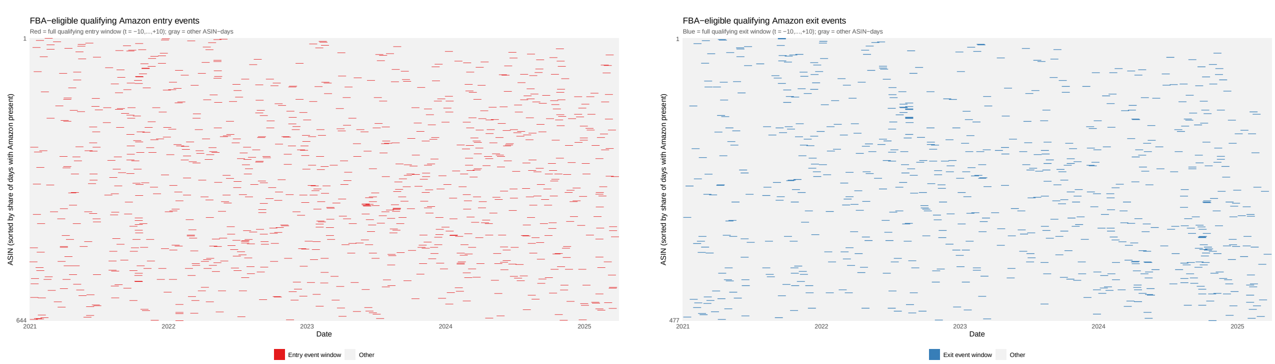

(a) Event windows around time of Amazon entry used in analyses.

(b) Event windows around time of Amazon exit used in analyses.

A.7.2 Additional Specifications

Tables 15 and 16 show the results from a more "naive", but more standard, two-way fixed effects (TWFE) specification using data at the daily level. We run separate regressions for each product

Table 13: Descriptives: Seller regimes (FBA as 3P price)

	Obs.	Share	N (ASINs)
Panel A: ASIN-day regime counts			
3P only	695,734	0.251	1,393
Amazon only	851,903	0.308	1,628
Both	316,128	0.114	1,425
Neither	906,321	0.327	1,779
Total	2,770,086	1.000	1,786

	Mean	Median	P10	P90
Panel B: Within-ASIN share of days in regime				
3P only	0.251	0.121	0.000	0.716
Amazon only	0.308	0.273	0.001	0.758
Both	0.114	0.047	0.000	0.324
Neither	0.327	0.237	0.029	0.755

	N (spells)	Mean duration	Median duration
Panel C: Regime persistence (spell duration in days)			
3P only	16,350	42.55	11.00
Amazon only	14,488	58.80	6.00
Both	13,165	24.01	4.00
Neither	18,104	50.06	6.00
Regime transitions per ASIN			
Mean	33.77		
Median	26.00		

Notes: Coffee category; ASIN-day panel (midnight prices), January 2021–March 2025. ASINs are restricted to those satisfying the event-study screen used in this script: at least 10 consecutive days with both Amazon and FBA non-missing, and at least 10 consecutive days with only the third-party price observed (Amazon missing). Regimes are defined by price availability of Amazon and FBA. "3P only": 3P price observed and Amazon missing. "Amazon only": Amazon observed and 3P missing. "Both": both observed. "Neither": neither observed.

Table 14: Descriptives on Amazon retail presence spell duration among switcher ASINs (FBA-active days)

	N	Mean	Std. Dev.	Median	P10	P90
Panel A: Spell-level distribution (spell duration in days)						
Amazon present	9,472	31.89	90.13	3.00	1.00	83.90
Amazon absent	9,722	66.23	158.36	9.00	1.00	176.00
Panel B: Per-ASIN number of spells						
Amazon present	1,261	7.51	11.41	3.00	1.00	17.00
Amazon absent	1,261	7.71	11.49	4.00	1.00	17.00
Panel C: Per-ASIN mean spell duration (days)						
Amazon present	1,261	65.14	92.35	34.00	4.21	157.50
Amazon absent	1,261	146.42	200.84	62.33	6.00	421.00

Notes: Coffee category; ASIN-day panel (midnight prices), January 2021–March 2025. ASINs are restricted to those satisfying the event-study screen: at least 10 consecutive days with both Amazon and NEW_third non-missing, and at least 10 consecutive days with NEW_third observed and Amazon missing. We further restrict to switcher ASINs in the 3P-active sample (days with FBA observed): at least one day where Amazon is present and one day where it is absent. A spell is a maximal run in which Amazon's presence status is unchanged. Panel A reports spell-duration distributions for Amazon-present and Amazon-absent spells. Panel B reports the per-ASIN number of spells (separately for Amazon-present and Amazon-absent spells). Panel C reports each ASIN's mean spell duration (separately for Amazon-present and Amazon-absent spells).

Figure 12: Event study of entry by Amazon as a seller into a given ASIN; daily Coffee data; split by Amazon's price level

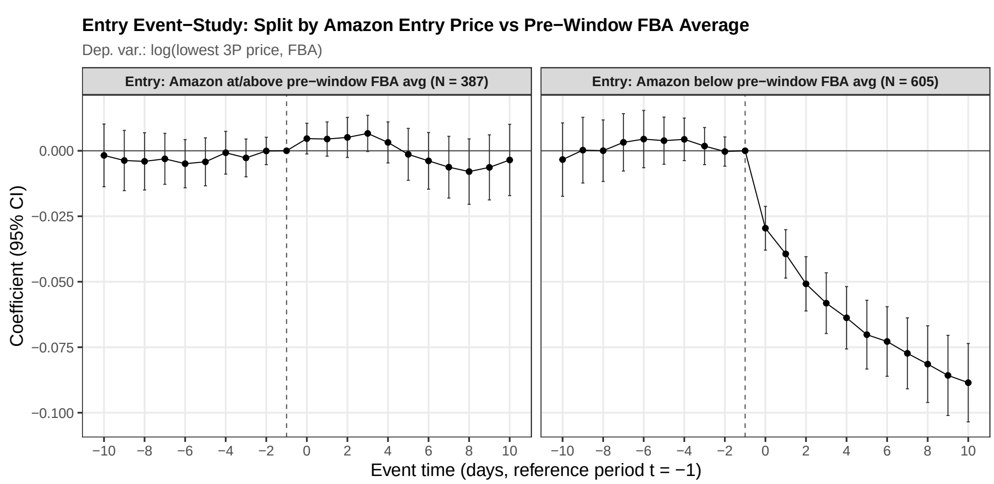

Figure 13: Event study of exit by Amazon as a seller into a given ASIN; daily Coffee data; split by Amazon's price level

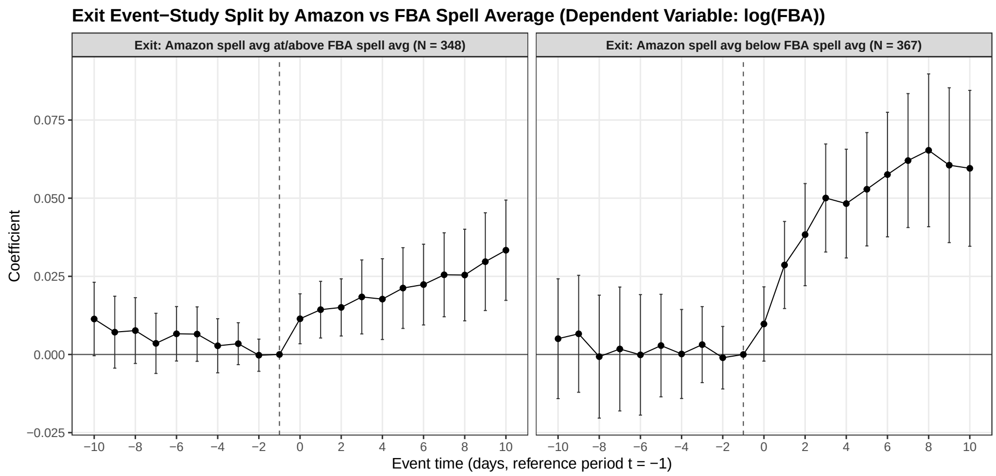

Figure 14: Event study of entry by Amazon as a seller into a given ASIN; daily Coffee data; split by pre-period amount of competition

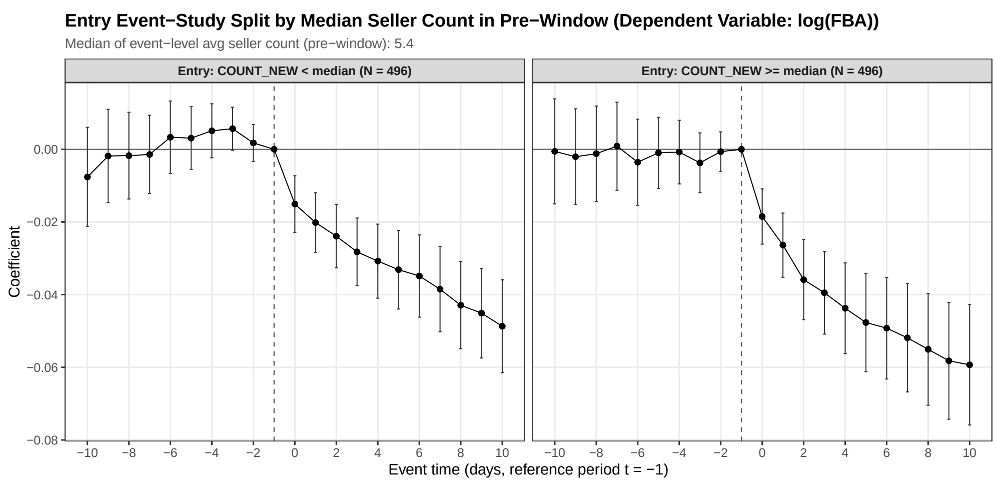

Figure 15: Event study of entry by Amazon as a seller into a given ASIN; daily Coffee data; split by post-period amount of competition

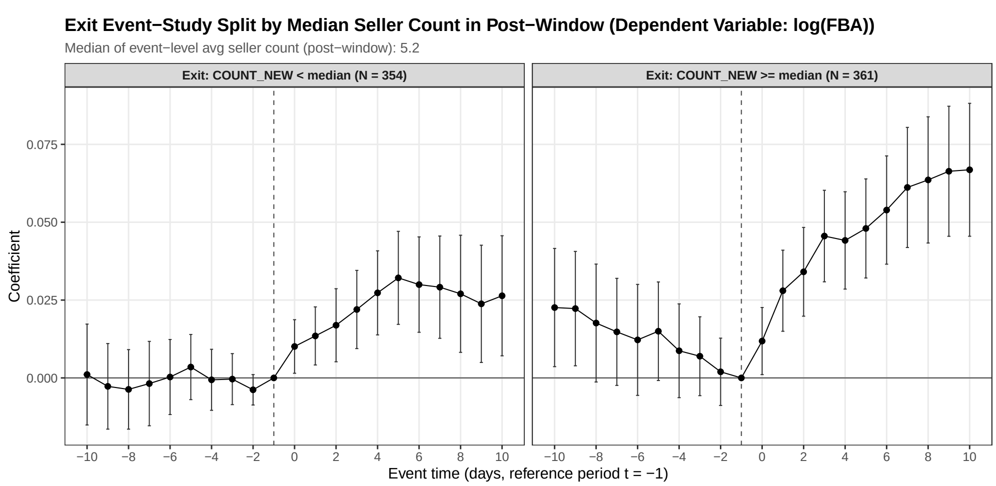

Figure 16: Event study of entry by Amazon as a seller into a given ASIN; daily Coffee data; longer event window horizon

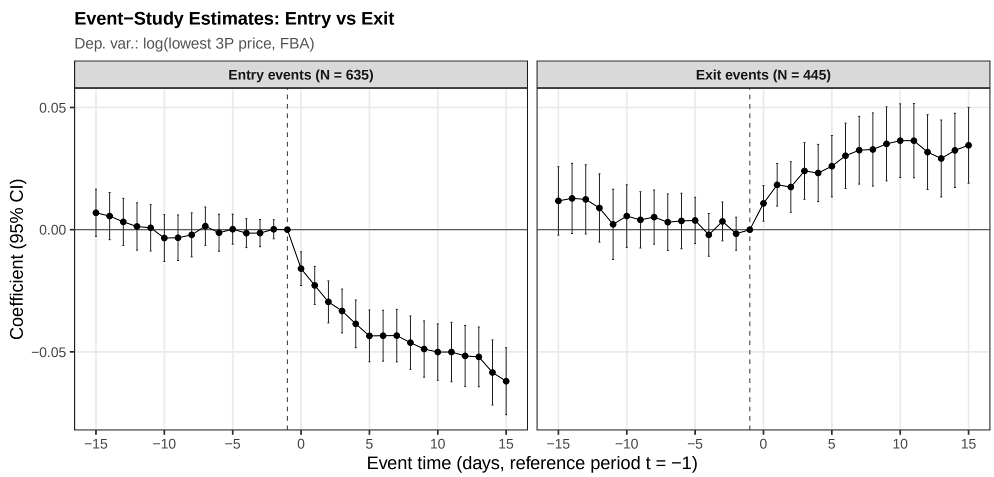

category. For product i on day t, we estimate:

$$\log FBA_{it} = \beta \cdot \mathbb{I}(\text{Amazon}_{it}) + \mu_i + \lambda_t + \varepsilon_{it} \tag{6}$$

where, as in the main text, FBA_{it} denotes the lowest price offered by third-party sellers under the FBA model of ASIN i on day t (FBA_{it} is defined analogously). $\mathbb{I}(\text{Amazon}_{it})$ is an indicator taking the value 1 if Amazon is active as a seller for product i at time t. The inclusion of product fixed effects (μ_i) and time fixed effects (λ_t) absorbs time-invariant product characteristics (e.g., quality) and common macro-shocks.

A.8 Pass-through: additional specifications

Tables 17, 18, and 19 results using alternative price measures, log specification, or an alternative subsample, respectively.

Table 15: Amazon's presence and 3P lowest price

Product	Coffee	Cocoa	Pasta	Shirts	TVs	USBhubs
Dependent Variable:	log(FBA)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
1(Amazon present)	-0.092***	-0.073***	-0.098***	-0.068***	-0.029***	-0.077***
	(0.006)	(0.020)	(0.010)	(0.010)	(0.008)	(0.020)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearmonthday	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	929,816	59,389	474,405	93,726	127,468	106,130
N ASINs	1,430	71	612	189	395	151
Adjusted R^2	0.90922	0.94750	0.92433	0.77291	0.98274	0.93941

Notes: Sample used: ASIN-day panels (midnight prices), January 2021–March 2025. In each category, ASINs are restricted to those satisfying the event-study screen used in this script: at least 10 consecutive days with both Amazon and FBA non-missing, and at least 10 consecutive days with only the third-party price observed (Amazon missing). We do not include SD cards or Washers due to a low number of qualifying ASINs. Clustered (asin) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Table 16: Amazon's presence and 3P lowest price

Product	Coffee	Cocoa	Pasta	Shirts	TVs	USBhubs
Dependent Variable:	log(FBM)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
1(Amazon present)	-0.030***	-0.053***	-0.055***	-0.018**	-0.046***	-0.067***
	(0.006)	(0.018)	(0.007)	(0.009)	(0.008)	(0.016)
Fixed-effects						
asin	Yes	Yes	Yes	Yes	Yes	Yes
yearmonthday	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	1,487,612	77,678	787,488	172,225	464,300	104,707
N ASINs	1,602	79	761	226	497	117
Adjusted R^2	0.91245	0.95494	0.91918	0.81099	0.96782	0.91219

Notes: Sample used: ASIN-day panels (midnight prices), January 2021–March 2025. In each category, ASINs are restricted to those satisfying the event-study screen used in this script: at least 10 consecutive days with both Amazon and FBA non-missing, and at least 10 consecutive days with only the third-party price observed (Amazon missing). We do not include SD cards or Washers due to a low number of qualifying ASINs. Clustered (asin) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Table 17: Pass-through: distributed lag regressions using Buy Box prices

Specification	K (lags)	Sum of Coefficients	Standard Error	T-statistic	Observations
Coffee:					
Buy Box Third-party	6	-0.18	0.32	-0.56	1279218
Buy Box Third-party	12	-0.37	0.49	-0.76	1279218
Buy Box Third-party	24	-0.49	0.66	-0.75	1279218
Buy Box Third-party	48	-0.18	0.66	-0.27	1279218
Buy Box Amazon	6	0.00	0.13	-0.04	275094
Buy Box Amazon	12	-0.15	0.18	-0.85	275094
Buy Box Amazon	24	-0.17	0.29	-0.61	275094
Buy Box Amazon	48	0.54	0.39	1.37	275094
Cocoa:					
Buy Box Third-party	6	0.13	0.05	2.42	128352
Buy Box Third-party	12	0.23	0.08	3.05	128352
Buy Box Third-party	24	0.18	0.13	1.39	128352
Buy Box Third-party	48	0.61	0.25	2.41	128352
Buy Box Amazon	6	0.03	0.10	0.26	9725
Buy Box Amazon	12	0.05	0.14	0.37	9725
Buy Box Amazon	24	-0.54	0.10	-5.27	9725
Buy Box Amazon	48	-1.08	0.38	-2.82	9725

Notes: Each row reports the cumulative pass-through estimate, $\sum_{k=0}^{K} \beta_k$, from the distributed-lag specification in equation 5, where K denotes the number of weekly lags of the commodity-price change (coffee and cocoa, respectively) included in the regression. Regressions are estimated separately for third-party seller prices conditionally on being shown in the Buy Box (*Buy Box Third-party*) and Amazon prices conditionally on being shown in the Buy Box (*Amazon*) using the baseline levels specification. The dependent variable is the weekly change in product price per gram, and the regressors are the contemporaneous and lagged weekly changes in the commodity price per gram. The sample is a weekly ASIN panel of products from January 2021 to March 2025. Standard errors are two-way cluster-robust. T-statistics are computed using these standard errors.

Table 18: Log pass-through: distributed lag regressions

Specification	K (lags)	Sum of Coefficients	Standard Error	T-statistic	Observations
Coffee:					
Third	6	0.02	0.01	2.16	1559229
Third	12	0.02	0.01	2.65	1559229
Third	24	0.04	0.01	3.31	1559229
Third	48	0.07	0.02	4.78	1559229
Amazon	6	-0.01	0.02	-0.65	288417
Amazon	12	-0.03	0.03	-1.09	288417
Amazon	24	-0.03	0.04	-0.81	288417
Amazon	48	0.02	0.04	0.54	288417
Cocoa:					
Third	6	0.01	0.01	1.47	149569
Third	12	0.03	0.01	2.70	149569
Third	24	0.05	0.02	2.87	149569
Third	48	0.11	0.02	4.70	149569
Amazon	6	0.02	0.02	0.73	11003
Amazon	12	0.02	0.04	0.49	11003
Amazon	24	-0.07	0.04	-1.86	11003
Amazon	48	-0.14	0.05	-3.09	11003

Notes: Each row reports the cumulative pass-through estimate, $\sum_{k=0}^{K} \beta_k$, from the distributed-lag specification in equation 5, where K denotes the number of weekly lags of the commodity-price change (coffee and cocoa, respectively) included in the regression. Regressions are estimated separately for logged third-party seller prices(*Third*) and logged Amazon prices (*Amazon*) using the baseline specification but in logs. The dependent variable is the weekly change in product price per gram, and the regressors are the contemporaneous and lagged weekly changes in the commodity price per gram. The sample is a weekly ASIN panel of products from January 2021 to March 2025. Standard errors are two-way cluster-robust. T-statistics are computed using these standard errors.

Table 19: Pass-through: distributed lag regressions, ASINs operating under dual mode only

Specification	K (lags)	Sum of Coefficients	Standard Error	T-statistic	Observations
Coffee:					
Third	6	0.33	0.17	1.95	237995
Third	12	0.55	0.23	2.37	237995
Third	24	0.69	0.28	2.44	237995
Third	48	0.92	0.56	1.65	237995
Amazon	6	-0.05	0.14	-0.34	238252
Amazon	12	-0.12	0.18	-0.66	238252
Amazon	24	-0.25	0.26	-0.95	238252
Amazon	48	0.25	0.36	0.70	238252
Cocoa:					
Third	6	0.46	0.19	2.38	10676
Third	12	0.37	0.21	1.77	10676
Third	24	-0.06	0.41	-0.13	10676
Third	48	1.49	0.74	2.02	10676
Amazon	6	0.28	0.36	0.78	7467
Amazon	12	0.35	0.58	0.61	7467
Amazon	24	-2.17	0.62	-3.52	7467
Amazon	48	-2.01	1.06	-1.89	7467

Notes: Each row reports the cumulative pass-through estimate, $\sum_{k=0}^{K} \beta_k$, from the distributed-lag specification in equation 5, where K denotes the number of weekly lags of the commodity-price change (coffee and cocoa, respectively) included in the regression. Regressions are estimated separately for third-party seller prices(*Third*) and Amazon prices (*Amazon*) using the baseline levels specification. The dependent variable is the weekly change in product price per gram, and the regressors are the contemporaneous and lagged weekly changes in the commodity price per gram. The sample is a weekly ASIN panel of products from January 2021 to March 2025; this time, we only include ASINs for which in at least 10% of ASIN-weeks with any offer, Amazon and third-party sellers are jointly offering the product. Standard errors are two-way cluster-robust. T-statistics are computed using these standard errors.