Demand Steering and Market Selection: Vertical Contracts and Algorithmic Allocation

Preliminary and Incomplete Draft

Muxin Li   Ksenia Shakhgildyan*

July 10, 2026

Abstract

We study how vertical agreements interact with platform algorithms to reshape competition in digital markets. We examine the 2018 Apple–Amazon brand-gating agreement, which restricted third-party sellers while granting Amazon privileged access to Apple products. Using cross-country variation and iPad product-level data, we document three findings. First, seller participation fell by nearly 90 percent and the number of products by more than half, driven primarily by the exit of older models, smaller but not less experienced or focused on iPad market sellers; yet seller exit alone caused no price increase and no quality improvement: prices on products not subsequently entered by Amazon actually declined despite reduced competition. Second, Amazon strategically entered newer, more popular, and higher-priced products, and introduced new models priced far above prevailing levels; its offers then captured the Buy Box over 70 percent of cases. Third, market power arises from the complementarity of both mechanisms: Amazon's algorithmic control sustained elevated prices on products it entered while steering buyer demand toward them, replacing price-based market selection with platform-controlled allocation.

Keywords: Digital Platforms, Brand Gating, Vertical Restraints, Algorithmic allocation, Exclusive Dealing.

JEL Codes: L42, D22, L51, L1, L2.

*Muxin Li, Fudan University, *muxin.econ@outlook.com*, Ksenia Shakhgildyan, Bocconi University and CEPR *ksenia.shakhgildyan@unibocconi.it*. We thank Egor Dolgov for research assistance. We thank the NET Institute `www.NETInst.org` for financial support.

1 Introduction

Vertical restraints and exclusive agreements have long been among central concerns of competition policy. On the one hand, such arrangements can generate efficiencies by addressing free-riding, protecting incentives to invest in brand quality, or reducing counterfeits. On the other hand, they can restrict market access, raise rivals' costs, and reinforce the market power of dominant firms. This trade-off has shaped antitrust debates for decades. In digital markets, the phenomenon becomes even more complex. Platforms such as Amazon are not merely intermediaries: they also compete as sellers and actively allocate demand through recommendation algorithms.

On Amazon, the Buy Box is a proprietary algorithm that selects, for each product listing, a single seller whose offer is prominently displayed with "Add to Cart" and "Buy Now" buttons. This featured slot captures the bulk of unit sales; a seller who does not win the Buy Box is effectively invisible to most consumers. Through this algorithm, Amazon exercises strong control over which seller captures consumer demand in its marketplace. As a result, vertical agreements can interact with platform governance to reshape not only competition among sellers, but also the consideration set of consumers.

This paper studies this interaction through the lens of the high-profile agreement between Amazon and Apple. On October 31, 2018, Amazon and Apple signed an agreement that took effect in January 2019 across major markets, including the U.S., U.K., France, Germany, Italy, Spain, Japan, and India. The agreement made Amazon an authorized reseller of nearly the full Apple and Beats product line while restricting the number of Apple-authorized resellers to no more than twenty per country. Hundreds of independent sellers were delisted overnight. Apple justified the deal as necessary to protect consumers and reduce counterfeits; Amazon secured a direct supply of Apple products while reducing competition from third-party resellers. The agreement is an ideal setting for empirical analysis: it operates at the product-category level, it applies simultaneously across multiple country marketplaces, and it has a sharp and well-documented implementation date.

This setting provides a unique opportunity to examine how vertical contracts and algorithmic allocation jointly determine market outcomes. In particular:

1 **Vertical Contract (Brand Gating):** restricts entry and reduces the set of active sellers.

2 **Algorithmic Allocation (Buy Box):** determines demand allocation across remaining sellers.

To analyze these effects, we assemble a new dataset that combines (i) Keepa marketplace data, covering daily observations of seller participation, product variety, Buy Box allocations, and prices across multiple countries, and (ii) the Amazon Reviews 2023 dataset, which includes consumer star ratings and review text. This dual approach allows us to capture both supply-side and demand-side dimensions of the market. Our identification strategy combines dynamic difference-in-differences and triple-differences designs, comparing the iPad market in treated versus control countries (Canada and Mexico) and contrasting iPads with their closest competitor, Samsung tablets. This framework enables us to isolate the causal effects of the agreement from global shocks and brand-specific trends.

Following the brand gating agreement, we observe several large and persistent changes in market outcomes. First, the agreement not only triggered the exit of independent Apple resellers but also significantly reduced number of available products: the average number of sellers per iPad fell by 90 percent, and the number of distinct models declined by more than half relative to the control group, driven primarily by the removal of older models from the platform. Product exit reflects cohort-based selection – products from earlier entry cohorts were disproportionately removed, with the difference between exiter and stayer entry-date distributions substantially stronger in treated than in control markets, ruling out natural product aging as the sole explanation.Second, Amazon's Buy Box share for Apple products surged from near-zero to over 55 percent, consistent with self-preferencing concernsnand Buy Box prices increased by 15–25 percent relative to controls, directly raising consumer expenditures.[1] Third, consumer ratings and sentiment analysis show no evidence of quality improvements or counterfeit removal: products that exited were indistinguishable from those that remained in terms of ratings, and average quality remained stable throughout. Strikingly, in treated markets the sellers expelled by brand gating were actually better-rated than those who remained, the reverse of what quality-based screening would predict and the opposite of the pattern observed in control markets.

We further document the profile of sellers who exited. They were not fringe operators: FBA adoption was nearly identical, months active on iPads were comparable, and iPad share of total storefront was similar. The sharpest contrast is seller scale: exiters averaged only 5 thousand total storefront ASINs versus 28 for stayers – a nearly six-fold difference – and listed 17 distinct iPad ASINs versus 38 for stayers. on average Critically, before the policy these smaller sellers were commercially dominant: exiters collectively held approximately 55–60% of iPad Buy Box share in treated markets, compared to 25–30% for stayers and Amazon combined. Their removal was not the elimination of a competitive fringe but the forced exit of the pre-policy Buy Box majority.

We further investigate the mechanisms underlying these market changes. We show that Amazon's behavior is highly strategic following the brand gating agreement: rather than entering the existing Apple catalog uniformly, it concentrated on newer, more popular, and higher-priced products, and simultaneously introduced a large number of new models at elevated prices immediately after the agreement. Once Amazon entered a product market, its offers captured the Buy Box in over 70 percent of cases, exceeding even the Buy Box win rate of third-party sellers who share Amazon's own fulfillment channel, which rules out logistics parity as the explanation.

We hightlight that seller exit alone surprisingly caused no price increase and no quality improvement. Rather, it is the complementarity between the vertical restraint and Amazon's platform strategy that drives the observed effects. Despite the sharp reduction in seller participation and competition, prices on products not subsequently entered by Amazon declined. Seller exit did, however, mechanically compress the range of prices offered for a given product, narrowing the competitive channel: the coefficient of variation of prices across sellers for the same ASIN fell from 11–15 percent pre-agreement to 4–7 percent post-agreement for stayers, while newly introduced products showed even less cross-seller dispersion at 2–4 percent. By contrast, products that Amazon chose to enter quickly accumulated reviews, improved ratings, and sustained high prices—outcomes driven by their privileged position in the Buy Box algorithm. This pattern reveals

[1] The Buy Box reallocation fell disproportionately on merchant-fulfilled resellers rather than those using Amazon's own fulfillment network, indicating the effect is not simply a byproduct of ordinary logistics-based ranking.

the key mechanism: *demand steering through the Buy Box transforms contractual access into market power, while the vertical restraint reinforces and sustains it.* Leveraging the brand-gating agreement, Amazon selectively entered newer, higher-priced, and more popular products; its algorithmic control then shielded these products from competitive price pressure. With the vertical restraint blocking new seller entry and older models removed from buyers' consideration sets, Amazon was able to maintain this advantage durably over time.

These findings carry direct implications for consumer welfare. The agreement raised effective prices in treated markets while eliminating more than half of the available product variety, concentrating the consideration set at higher price points. The affordable tier of the market—previously served by independent resellers offering older models at competitive prices—was largely displaced, with no offsetting improvement in average product quality. Beyond prices and variety, Amazon's Buy Box dominance generated an informational distortion: by channeling the bulk of transactions to its own listings, it allowed those products to accumulate reviews rapidly, further entrenching its algorithmic advantage. Together, these effects indicate that the agreement imposed significant costs on consumers through the joint operation of contractual foreclosure and algorithmic demand steering.

Related Literature Our study contributes to three strands of the literature. The first strand is the literature on hybrid platforms—intermediaries that both host third-party sellers and compete with them directly. This dual role creates inherent conflicts of interest, raising concerns about steering and discrimination [Hagiu and Wright, 2020, Motta, 2023]. Theoretical work shows that vertical integration generates incentives for algorithmic bias, as hybrid platforms tilt search results, rankings, or default recommendations toward their own offers [De Corniere and Taylor, 2014, 2019, Inderst and Valletti, 2011, Kittaka and Sato, 2022, Padilla et al., 2022, Raskovich, 2007, Zennyo, 2022].

More broadly, hybrid platforms distort competition through governance policies—such as restricting seller participation, biasing innovation, or deploying manipulative design features [Karle et al., 2020, Teh, 2022]—or by moderating content and delisting sellers [Bedre-Defolie et al., 2024, Casner, 2020, Jiménez Durán et al., 2022, Liu et al., 2022]. On the empirical side, self-preferencing at Amazon has been documented in search rankings [Farronato et al., 2023] and Buy Box assignment [Raval, 2022]; related work studies Amazon's competitive strategies and marketplace dynamics [Chen and Tsai, 2024, Gómez-Losada and Duch-Brown, 2019], its selective entry into product categories [Zhu and Liu, 2018], the welfare effects of its review system [Reimers and Waldfogel, 2021], and the regulation and welfare consequences of platform self-preferencing more broadly [Decarolis and Li, 2025, Decarolis et al., 2025, Gutiérrez, 2022, Waldfogel, 2024]. Our paper contributes to this strand in three respects. First, we document a mechanism that is qualitatively distinct: Amazon's self-preferencing was enabled not through internal platform policies but through a bilateral vertical contract with a manufacturer, which restructured the seller landscape and allowed the Buy Box algorithm to steer demand toward Amazon's listings. Second, we show that Amazon's entry decisions are highly strategic—concentrating on newer, more popular, and higher-priced products—and that its Buy Box control sustains prices and captures demand in ways competing sellers cannot replicate even under identical contractual conditions. Third, by decomposing Buy Box outcomes by fulfillment channel, we isolate algorithmic self-preferencing from a purely logistics-based explanation—a distinction that prior work relying on aggregate Buy Box win rates has not been able to draw.

The second branch regards vertical restraints and exclusive dealing. Classical work establishes how vertical agreements can foreclose rivals or soften competition [Aghion and Bolton, 1987, Rey and Tirole, 1986], while more recent contributions formalize conditions under which restraints generate efficiencies versus exclusionary risks [Drugov and Jeon, 2019, Hagiu et al., 2022, Rey and Vergé, 2004]. The Apple–Amazon agreement closely resembles exclusive dealing: it restricts distribution to Amazon and a narrow set of authorized resellers, eliminating intra-brand competition. We extend this literature in two ways. First, we provide cross-country empirical evidence on the magnitude of foreclosure effects across multiple markets simultaneously. Second, we show that the vertical restraint alone is insufficient to generate sustained price increases—its anticompetitive effects materialize only through interaction with the platform's algorithmic demand allocation, a complementarity that existing theory has not formalized.

The third branch focuses on brand protection, counterfeits, and quality assurance. Prior work shows that unauthorized sellers can erode consumer trust and undermine brand investment [Elfenbein et al., 2015, Hui et al., 2016], while certification intermediaries use screening mechanisms to enhance market quality [Lizzeri, 1999]. A related literature examines platform liability and how holding intermediaries responsible for harmful products shapes incentives for quality, innovation, and seller participation [Daughety and Reinganum, 1995, 2006, De Chiara et al., 2024, José Ganuza et al., 2016, Landes and Lichtman, 2003, Polinsky and Shavell, 2012, Zennyo, 2023]. These studies suggest that brand gating may be efficiency-enhancing if it removes low-quality or counterfeit products. Our findings challenge this view directly: products that exited the market were statistically indistinguishable from those that remained in terms of ratings and review sentiment, and expelled sellers were rated no worse—if anything, slightly better—than those who stayed. We thus provide novel empirical evidence that the quality-based justification offered by Apple and Amazon is not supported by the data.

The remainder of the paper is organized as follows. Section 2 provides institutional background on Amazon, Apple, and the brand-gating agreement. In Section 3 and Section 4, we describes the data and econometric methodology. We then document in Section 5 the agreement's effects on market outcomes and investigate underlying mechanisms in Section 6. In Section 7 we assess the welfare implications and conclude in Section 8.

2 Institutional Background

Digital platforms have become central to modern commerce, especially in markets shaped by network effects and two-sided dynamics. In such environments, a platform's value grows as more users participate on both sides of the market. This is the essence of a two-sided market, where platforms serve as intermediaries between distinct groups (e.g., consumers and sellers), and the interaction between these groups generates positive cross-group network effects. More buyers attract more sellers, and more sellers, in turn, attract more buyers. This positive feedback loop often produces "winner-takes-all" outcomes, where a few dominant platforms capture most of the market and create entry barriers for competitors. Two firms that shape the digital economy are Amazon and Apple. As of 2025, Apple ranks as the world's most valuable brand and Amazon the fourth, with brand values of $574.5 billion and $356.4 billion, respectively.[2]

Amazon, founded in 1994 as an online bookstore, has evolved into the world's largest e-commerce platform. By 2025, it hosted around 3.6 million active sellers and offered more than 350 million products.[3] The platform attracts extraordinary consumer traffic, with an estimated 2.8 billion monthly visits.[4] In the United States, Amazon accounts for about 38 percent of online retail sales-over six times that of its closest competitor, Walmart, which holds about 6 percent.[5] In some markets, its dominance is even starker: in Italy, Amazon captures about 75 percent of marketplace visits,[6]and in certain European countries, its share of online consumer electronics sales approaches 90 percent.[7] Amazon operates under a hybrid model: both as a marketplace and as a retailer. As a marketplace, it provides millions of third-party merchants with access to its consumer base in exchange for commissions, advertising, and fulfillment fees. As a retailer, it purchases goods wholesale and sells them directly. This dual role enhances Amazon's efficiency but blurs the line between neutral intermediation and direct competition. Amazon not only hosts sellers but also competes with them, with the power to influence rankings and product visibility–an arrangement that has drawn persistent antitrust scrutiny.

Apple, founded in 1976, has built market power through a different strategy: vertical integration. Rather than simply intermediating, Apple designs hardware, software, and services as a tightly integrated ecosystem. This model fosters brand loyalty, raises switching costs, and sustains high margins in premium segments. In 2023, Apple surpassed Samsung to claim the largest global smartphone share for the first time in over a decade, accounting for around 20 percent of shipments.[8] In the tablet market, it maintained over 48 percent of worldwide share, while in audio it led globally with 81.8 million headphones shipped.[9] These figures underscore Apple's strength in premium markets and its ability to exercise strict control over distribution.

The relationship between Apple and Amazon is complex, characterized by both complementarity and rivalry. Apple depends on Amazon to reach the world's largest online customer base, yet Amazon markets its own hardware, such as Fire Tablets and Echo Buds, that directly compete with iPads and AirPods.

This tension subtly changed with the Global Tenants and Authorized Reseller Agreement signed on October 31, 2018. This landmark deal granted Amazon access to almost the entire Apple and Beats products portfolio, establishing a direct supply from Apple.[10] Prior to the deal, Apple products on Amazon were primarily sold by independent resellers. The agreement fundamentally restructured this market: on Amazon sites in the U.S., U.K., France, Germany, Italy, Spain, Japan, and India, access to selling Apple and Beats products was restricted to a closed, country-specific list of named resellers, capped at no more than a fixed number per country—down from hundreds previously. Independent resellers, by contrast, were delisted as of January 4, 2019, under the so-called *Brand Gating Clause*, with Amazon contractually required to block any seller

[2]See DemandSage, "94 Amazon Statistics & Market Share 2026".
[3]See Capital One Shopping, "Number of Products on Amazon".
[4]See AMZScout, "Amazon Statistics".
[5]See DemandSage, "94 Amazon Statistics & Market Share 2026".
[6]See Statista, "Most popular online marketplaces in Italy in December 2023, based on share of visits".
[7]See Statista, "Share of iPhone unit shipments in smartphone shipments worldwide from 3rd quarter 2007 to 1st quarter 2026".
[8]See "Apple #1 In Global Smartphone Market For First Time Ever".
[9]See Statcounter, "Tablet Vendor Market Share Worldwide" and Tadvisor, "Headphones (Global Market)".
[10]Apple's HomePod smart speaker was not made available on Amazon.

not appearing on Apple's list.[11]

To obtain "Authorized Reseller" status, retailers were required to sign a dedicated agreement with Apple, which offered discounts and rebates in exchange for strict compliance with Apple's resale conditions.[12] However, authorized reseller status was a necessary but not sufficient condition for platform access: sellers who already held Apple Authorized Reseller or Apple Premium Reseller status could nonetheless be excluded if they did not appear on the named list. The mechanism thus operated as a double filter—requiring both official authorization and explicit inclusion—so that authorization did not guarantee access. Modifications to the list required Apple's approval, and Apple retained full discretion to add, remove, or replace resellers based on its own internal considerations.

Beyond the general terms governing Authorized Reseller status, the agreement also restructured Apple's direct commercial relationship with Amazon. The complaint underlying the parallel U.S. antitrust litigation alleges that Apple agreed to supply Amazon with consistent inventory at a discount of up to 10 percent, contingent on Apple's ability to keep the number of resellers on its marketplace low, tying Apple's own margin on Amazon sales directly to the scale of the reseller cap it imposed.[U.S. federal court antitrust complaint concerning alleged Apple–Amazon agreement, 2023] This detail clarifies the economic structure of Amazon's role in the agreement: rather than earning a referral commission on third-party sales in the conventional marketplace sense, Amazon purchased Apple and Beats inventory outright and resold it at retail, retaining the markup as its own margin. Consistent with this wholesale structure, Apple's own resale terms specify that title and risk of loss for Apple products pass to the reseller upon shipment from Apple's shipping location, so that Amazon owned the inventory it sold rather than acting as a neutral intermediary for Apple-fulfilled listings.[13]

Apple justified the agreement as a means to improve brand control and reduce counterfeiting. However,regulatory authorities suggest that inclusion was not governed by objective or qualitative criteria. Italy's AGCM found that the numerical cap was a global limit proposed by Apple, that sellers were selected individually without reference to performance standards, and that Apple had internally resisted incorporating any such criteria into the contracts—leading the authority to characterize the mechanism not as a genuine selective distribution system but as a controlled restriction on market entry.[14] Amazon, meanwhile, secured privileged access to Apple's products while reducing intra-brand price competition from resellers. For consumers, the effects were likely mixed: authenticity and consistency might have been improved, but discount-driven rivalry declined, potentially also softening inter-brand competition with rivals like Samsung.

Unsurprisingly, the agreement has drawn significant regulatory attention. In Italy, the competition authority fined Apple €134.5 million and Amazon €68.7 million in November 2021 for anticompetitive coordination in Apple and Beats distribution [Autorità Garante della Concorrenza e del Mercato (AGCM), 2023], though the €203 million penalty was annulled in 2022 on

[11] Starting on 16 November 2021 (after the initiation of investigations), the list was expanded to include Apple Premium Resellers, Authorized Apple Resellers, and Retailers (regardless of the market), but non-authorized resellers continued to be excluded. See CNBC, "More Apple products are coming to Amazon under a new deal".

[12] Authorized Resellers are offered discounts and rebates in exchange for adherence to Apple's terms and conditions governing resale. See US "Amended Class Action Complaint".

[13] See Apple, "Apple Store for Resellers Terms and Conditions".

[14] See Autorità Garante della Concorrenza e del Mercato (AGCM) [2023].

procedural grounds, leaving substantive questions unresolved.[15] In Spain, the CNMC imposed a record-breaking €194 million fine in July 2023–€143.6 million on Apple and €50.5 million on Amazon–for clauses restricting resellers, advertising, and targeted marketing.[Comisión Nacional de los Mercados y la Competencia (CNMC), 2023] Spain's high court suspended enforcement pending appeal in early 2024.[16] Germany's Federal Cartel Office opened an investigation in 2020 into brand-gating as potential abuse of dominance, and in the United Kingdom, a proposed £494 million collective action alleged collusion to foreclose resellers and inflate prices. Though dismissed in 2025 on procedural grounds, the case reflected growing concern about such vertical arrangements.[17] At the EU level, both firms have been formally designated as "gatekeepers" under the Digital Markets Act, which imposes obligations to preserve platform neutrality and avoid exclusionary conduct.[18] In the U.S., a parallel antitrust litigation has alleged unlawful horizontal collusion and group boycott under the Sherman Act [U.S. federal court antitrust complaint concerning alleged Apple–Amazon agreement, 2023].

3 Data

Keepa Data Our analysis relies primarily on data obtained from Keepa, a third-party API that provides comprehensive and granular information on Amazon's marketplace.[19] Keepa has systematically monitored Amazon platforms in Canada, France, Germany, Italy, Japan, Mexico, Spain, and the United States. Products are uniquely identified by their Amazon Standard Identification Number (ASIN), which ensures consistency and comparability across markets and product categories. In this paper, we focus on the period from May 2018 to September 2019, encompassing several months before and after the introduction and early implementation of the Apple–Amazon agreement. This time frame allows us to observe the state of the marketplace before the agreement came into force, as well as the market changes that followed.

To construct the dataset, we begin by identifying ASINs corresponding to Apple and Samsung tablets, which together represent a dominant share of the tablet market.[20] We then group these ASINs into product markets and extract product-level information at specific points in time, creating a structured dataset that captures competition on the supply side as well as its impact on consumers.

One of the most valuable features of the Keepa data is that it allows us to track product and seller participation over time. For each ASIN, we identify through a variety of measures the first and last dates of availability on Amazon. As a consequence, we are able to analyze product entry and exit. The data also records the set of sellers associated with each product at a given moment. This means we can see not only whether a tablet model remains available, but also how the pool of sellers changes over time. This dimension is especially important when evaluating an agreement

[15] See AGCM, Press Release, December 2021 and StartMag, "Why the TAR Rejected the AGCM Fine to Amazon and Apple".

[16] See CNMC, Press Release, July 2023 and Yahoo Finance, "Spain's High Court Suspends €209 Million Apple–Amazon Fines on Apple, Amazon amid appeal".

[17] See Reuters, "Apple, Amazon Fight Off £600 Million UK Lawsuit Over Alleged Collusion".

[18] See European Commission, Press Release on Digital Markets Act Gatekeeper Designations.

[19] See Keepa.

[20] Globally in 2024, Apple's iPads had around 55% share of tablet shipments, while Samsung's share was about 27%. In the U.S. market in 2024, Apple's share was around 57%, while Samsung's was around 16%. See BankMyCell, "Tablet Market Share Globally & US (2026)".

that restricts seller participation. Prior to the agreement, a single model of an Apple iPad might be offered by dozens of merchants competing on price and fulfillment. Following the agreement, however, only Amazon itself and a small set of authorized sellers were permitted to sell Apple products. The ability to observe the number of sellers as well as seller identities at the product level is, therefore, central to measuring how the agreement reshaped the structure of competition.

The data also provides visibility into the Buy Box algorithm. On a daily basis, we are able to observe the identity of the seller that won the Buy Box. This makes it possible to study the algorithm that determines which offer is promoted prominently to consumers, and Amazon's Buy Box share over time. If Amazon's own offers were disproportionately likely to win the Buy Box despite being priced higher than competing offers, this would suggest that the algorithm gave Amazon a structural advantage. Instead, if Buy Box outcomes aligned with observable characteristics such as price, fulfillment by Amazon, or seller ratings, this would indicate that the algorithm operated in a neutral or efficiency-enhancing way. Observing these patterns before and after the agreement provides a direct test of whether Amazon altered the Buy Box algorithm to strengthen its new role as an authorized Apple reseller.

Another strength of the Keepa data is its ability to track prices over time. Most importantly, we observe the Buy Box price. The Buy Box seller is displayed prominently to consumers and drives the majority of transactions. Given that, we can observe the price consumers are most likely to pay. That allows us to study whether the Buy Box algorithm steers demand toward higher-priced offers. In the future analysis, we examine whether the agreement reduced price competition for Apple tablets and Samsung tablets. Additionally, we can test whether Amazon's own pricing behavior changed once it became an authorized Apple reseller.

In addition to seller participation, pricing, and Buy Box assignment, the Keepa dataset also records consumer feedback. In particular, we use the average product rating and seller rating as a proxy for product quality. It allows us to examine whether the lower quality products or sellers were expelled from the Amazon marketplace, which would be an indication of the exit of the counterfeit products. This is critical for assessing the welfare implications of the agreement. If instead there was no statistically significant difference between the products or sellers that stayed on Amazon and the ones that exited, the agreement may have reduced consumer welfare by decreasing the choice of the highly rated products. Conversely, if lower-rated products or sellers exited disproportionately, the agreement may have raised average quality by filtering out weaker offerings. Linking changes in market structure to this consumer-facing indicator helps us connect competition outcomes to broader measures of welfare.

In Table 1, we report summary statistics for iPad listings before and after the Apple–Amazon agreement, separately for control and treated countries between May 2018 and May 2019. To capture key patterns in the data, we include several variables. *Sellers* is the monthly count of distinct seller IDs active on Apple tablet listings. *Seller / ASIN* is the monthly average number of distinct sellers offering each ASIN on a given day. *Product Number* refers to the monthly average number of distinct ASINs. The *Buy Box Price* is the monthly average Buy Box price in euros, adjusted for inflation. As the majority of transactions are through Buy Box recommendations, we view the Buy Box price as a critical measurement to shed light on the welfare implications. The *Rating* is the monthly average product rating and represents the quality measure. *Amazon BB Share* is the Amazon's share as the Buy Box seller.

Table 1: Summary Statistics of iPads Pre- and Post-Agreement

Country	Group	Period	Sellers	Sellers / ASIN	Product Number	Buy Box Price	Rating	Amazon BB Share
CA	Control	Before	55	1.39	312.67	563.02	4.33	0.01
CA	Control	After	56	1.67	379.00	516.55	4.21	0.03
MX	Control	Before	24	1.38	186.00	663.08	4.47	0.03
MX	Control	After	20	1.14	167.00	737.09	4.25	0.01
DE	Treated	Before	350	5.66	619.50	700.46	4.09	0.01
DE	Treated	After	21	1.94	269.50	758.46	4.09	0.40
ES	Treated	Before	185	3.77	542.17	683.80	4.13	0.06
ES	Treated	After	8	1.83	195.75	845.16	4.08	0.59
FR	Treated	Before	165	4.06	417.50	724.62	3.87	0.03
FR	Treated	After	11	1.21	190.50	850.12	3.97	0.55
IT	Treated	Before	194	4.57	401.50	623.61	4.26	0.01
IT	Treated	After	14	1.65	211.25	853.92	4.27	0.56
JP	Treated	Before	448	3.42	208.33	511.96	4.51	0.00
JP	Treated	After	14	1.18	53.50	951.44	4.32	0.86
US	Treated	Before	868	7.08	451.67	467.42	4.33	0.00
US	Treated	After	80	3.57	391.25	475.46	4.26	0.21

Notes: *Before* period: May 2018 – October 2018. *After* period: February 2019 – May 2019. The January–December 2018 transition months are excluded from the *After* period. For all variables, monthly averages across active ASINs were first computed, then averaged across months within each period. *Sellers* is the average monthly count of distinct seller IDs active on Apple tablet listings. *Sellers/ ASIN* is the average number of daily new-condition sellers per ASIN. *Product Number* is the average number of active ASINs per month. *Buy Box Price* is the Buy Box price (including shipping) in euros, adjusted for inflation. *Rating* is the average product rating. *Amazon BB Share* is the Amazon's share as the Buy Box seller, based on the seller-ID history.

In the control group, the number of sellers and products remained relatively stable, with only modest variations across periods. By contrast, the treated markets experienced a sharp decline in seller participation after the agreement. For example, in France the average monthly number of sellers decreased from 165 to just 11. Other countries experienced even larger drops. Notably, the decline in seller counts was accompanied by a drop in product variety, as measured by the number of distinct ASINs. At the same time, average Buy Box prices increased notably after the agreement, while in the control countries the pattern is relatively stable. Moreover, Amazon share as a Buy Box seller increase dramatically in the treated countries, while it remained close to zero in control countries. The table, therefore, highlights the central empirical pattern: treated countries saw a significant contraction in seller count and product variety after the agreement, accompanied by upward pressure on prices, while control markets displayed much milder changes over the same period. Moreover, product rating remained relatively stable.

Amazon Reviews Data The Amazon Reviews 2023 dataset constitutes one of the most comprehensive publicly available collections of customer reviews in the United States.[21] While the dataset covers a wide range of product categories, our analysis focuses on reviews of iPads. Each observation corresponds to an individual review and contains the product's ASIN, the review date, the star rating, the review title and text, the number of helpfulness votes, and an indicator of whether the purchase was verified on Amazon. These data allow us to complement our marketplace analysis with an additional consumer perspective. Specifically, we compare ratings for products that remained in the market with those that exited, and we implement sentiment analysis to construct an alternative measure of consumer sentiment from review text. The data further enable us to examine whether reviews became more favorable or critical following the Agreement, whether consumers explicitly noted changes in availability or pricing, and whether the share of verified

[21] See Amazon Reviews 2023 Dataset.

purchases shifted once unauthorized sellers were excluded.

In sum, the review data provide an important demand-side dimension, capturing how consumers perceived the products and sellers available to them. Whereas Keepa data allow us to study structural outcomes on the supply side—such as seller participation, pricing dynamics, and Buy Box allocation–the review data shed light on consumer responses in the form of ratings, sentiments, and narratives. In summary, the two datasets provide a more complete picture of the Apple–Amazon agreement's impact on competition and consumer welfare. The main limitation is that the Amazon Reviews 2023 dataset is restricted to the U.S. market.

Additional Data As complementary sources, we also use monthly currency exchange rate data from the ECB Data Portal[22] and country-level inflation data from the World Bank[23], which allow us to translate prices in euros and adjust according to inflation rates.

4 Empirical Methodology

We study the causal effect of the Apple–Amazon agreement on the seller participation, product availability, prices, and quality using two complementary strategies: (i) a panel difference-in-differences (DiD) event study for Apple products, and (ii) a dynamic triple-differences (DDD) specification contrasting Apple with Samsung. Both designs are estimated on an ASIN–country–month panel.

4.1 Panel DiD Event Study

We let Y_{ict} denote the outcome of interest for ASIN i in country c and month t, $\mathit{Treated}_c \in \{0, 1\}$ denote an indicator for countries covered by the agreement, and t_0 denote the first month of the Agreement (November 2018).[24] We estimate the following dynamic specification:

$$Y_{ict} \;=\; \alpha_c \;+\; \gamma_t \;+\; \sum_{k \neq -1} \beta_k \left[\mathit{Treated}_c \times 1\{t - t_0 = k\} \right] \;+\; \varepsilon_{ict}, \tag{1}$$

where α_c measures country fixed effects and γ_t measures the monthly fixed effects. The omitted category is $k = -1$ (the month before t_0) as the baseline. The coefficients β_k trace the dynamic treatment effects relative to the baseline, enabling inspection of pre-trends ($k < 0$) and the post-implementation path ($k \geq 0$). Figure A.1 provides stylized schematics of the coefficients estimated by (1), with standard errors to be clustered at the country level.

4.2 Dynamic Triple-Differences (DDD)

The event study compares treated to control countries for Apple. To adress the remaining platform-level shocks common to all brands, we compare Apple with an untreated close competitor

[22] See ECB Data Portal, Exchange Rate Data.
[23] See World Bank Open Data.
[24] In robustness checks, we also consider January 2019, when Brand Gating was fully implemented.

(Samsung) and estimate a DDD model:

$$Y_{ict} \;=\; \alpha_{cb} \;+\; \gamma_t \;+\; \sum_{k\neq -1} \theta_k \left[\mathit{Treated}_c \times \mathit{Apple}_b \times 1\{t - t_0 = k\} \right] \;+\; \varepsilon_{ict}, \tag{2}$$

where Country×brand fixed effects α_{cb} absorb time-invariant brand differences within each country; γ_t captures global monthly shocks. The dynamic coefficients θ_k measure Apple-specific effects in treated countries at each event time k, relative to Samsung and to control markets. Figure A.2 provides stylized schematics of the coefficients estimated by (2). The shapes are illustrative; empirical figures in the results section use the estimated coefficients and their confidence intervals.

Identification requires parallel trends in the *Apple–Samsung difference* between treated and control countries in the absence of the agreement. Inspection of the pre-period coefficients $\{\theta_k : k < 0\}$ provides an internal check. The dynamic path $\{\theta_k : k \geq 0\}$ documents the timing and persistence of Apple-specific effects.

4.3 Discussion of Assumptions and Diagnostics

The event-study design requires parallel trends between treated and control countries for Apple in the pre-period; the DDD requires parallel trends in the *Apple–Samsung difference* across treated and control units. We display full event-time paths with confidence intervals.

5 Market Changes

5.1 Seller Exit and Product Variety

We begin by investigating the most direct channel through which the Apple–Amazon agreement reshaped the marketplace: the forced exit of unauthorized third-party resellers and the resulting change in the products available in consumers' consideration sets.

Seller exit. To begin this analysis, we examine the monthly averages of the number of sellers (including both third-party resellers and Amazon) offering iPads in treated and control markets, as shown in Figure 1a. Before the agreement, treated countries averaged more than 350 unique monthly sellers compared with less than 50 in the control group. After the agreement took effect, seller counts in treated markets fell sharply and stabilized at less than 20, while the control group remained unchanged.

In Figures 1b and 1d, we report the corresponding event-study and dynamic DDD estimates, which confirm this pattern. Pre-period coefficients are centered at zero and generally statistically indistinguishable from zero in both specifications, supporting the parallel pre-trends assumption. Post-implementation, the event-study coefficients turn sharply negative, stabilizing near −2.7 log points through months three to six post-agreement implying a roughly **93%** decline in active Apple resellers. The triple-difference estimates, which use Samsung tablets as a within-country benchmark, yield a consistent result. The Apple–Samsung seller gap widens abruptly in treated markets after February 2019 and remains flat in the control group (Figures A.3d and A.3c), which rules out platform-wide trends or common seasonal variation. The regional breakdown in Figure A.3e also shows that the decline is not concentrated in any single country but appears consistently across all treated markets.

Figure 1: Impact of the Apple–Amazon agreement on the number of Apple tablet sellers on Amazon

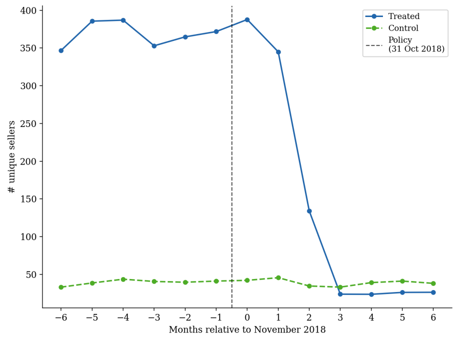

(a) Raw averages: treated vs control

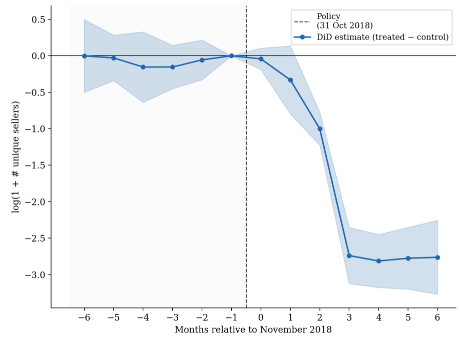

(b) Event Study: iPad sellers

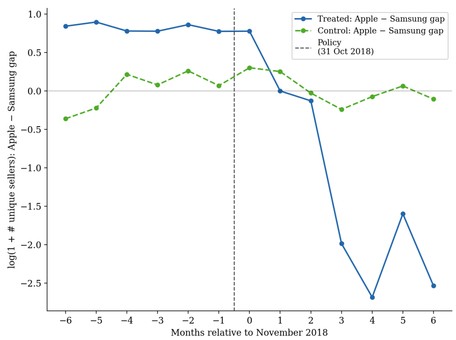

(c) Diagnostic: Apple–Samsung difference

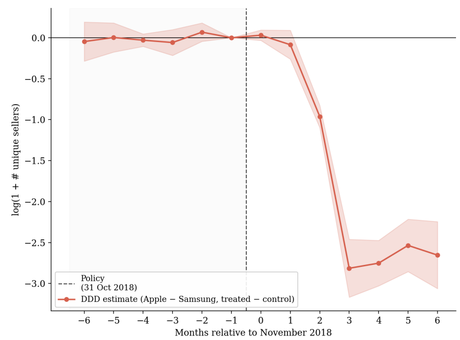

(d) Dynamic Triple-Difference: Apple vs Samsung

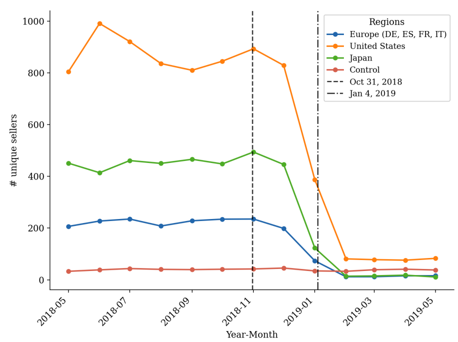

(e) Average monthly number of iPad sellers by region

Note: Figure A.3a reports raw averages; Figure A.3b shows the event-study estimates; Figure A.3c shows the raw diagnostic Apple–Samsung difference; Figure A.3d presents the dynamic triple-difference estimates contrasting Apple with Samsung; and Figure A.3e shows the raw averages by region.

Product variety. Notably, the reduction in the number of sellers also led to a decline in the number of products offered on the platform. As shown in Figure 2a, before the agreement, treated countries averaged about 450 unique Apple tablet models listed per month, compared with approximately 250 in the control group. After the agreement took effect, ASIN counts in treated markets fell sharply and stabilized at under 200 by early 2019, while control markets remained stable or continued to rise. This decline is specific to Apple: as shown in Figure 2c, no comparable reduction is observed for Samsung tablets over the same period.

In Figures 2b and 2d, we report the corresponding event-study and dynamic DDD estimates, which confirm this pattern. Pre-period coefficients are centered at zero and generally statistically indistinguishable from zero in both specifications, supporting the parallel pre-trends assumption. Both results suggest a sharp reduction of around **55%** in active iPad listings relative to baseline. The regional breakdown is provided in Figure 2e.

Figure 2: Impact of the Apple–Amazon agreement on the number of Apple tablets listed on Amazon

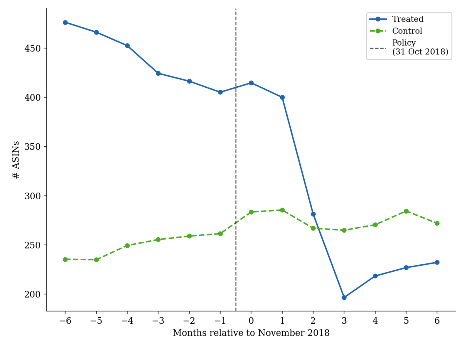

(a) Raw averages: treated vs control

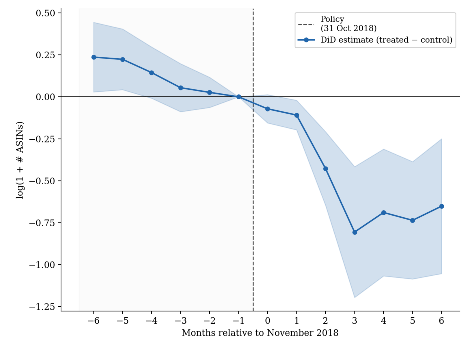

(b) Event Study: iPad listings

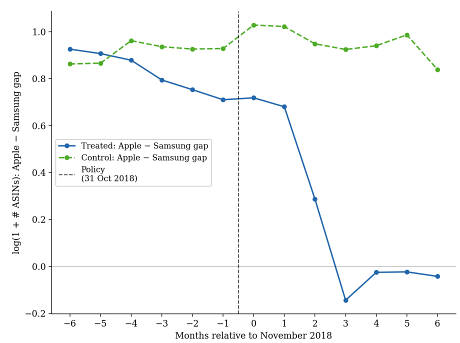

(c) Diagnostic: Apple–Samsung difference

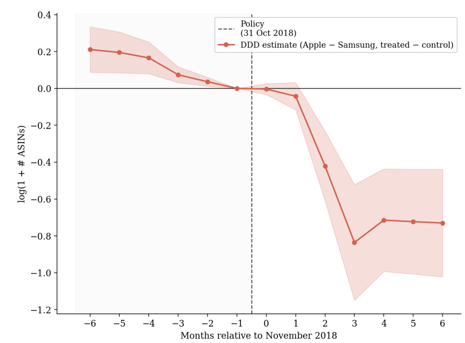

(d) Dynamic Triple-Difference: Apple vs Samsung

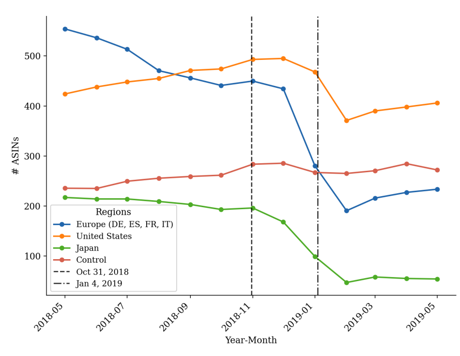

(e) Average monthly number of Apple tablets (unique ASINs) by region

Note: Figure 2a reports raw averages; Figure 2b shows the event-study estimates; Figure 2c shows the raw diagnostic Apple–Samsung difference; Figure 2d presents the dynamic triple-difference estimates contrasting Apple with Samsung; Figure 2e shows the shares by region.

Compositional shift toward newer products. Beyond the reduction in the number of listings, the agreement also shifted the composition of available products toward more recently introduced models. Figure 3 plots histograms of product introduction dates in treated and control markets for the pre-agreement period (May–October 2018) and the post-agreement period (January–September 2019).

Figure 3: Histogram of product introduction dates: treated vs. control markets, before and after the agreement

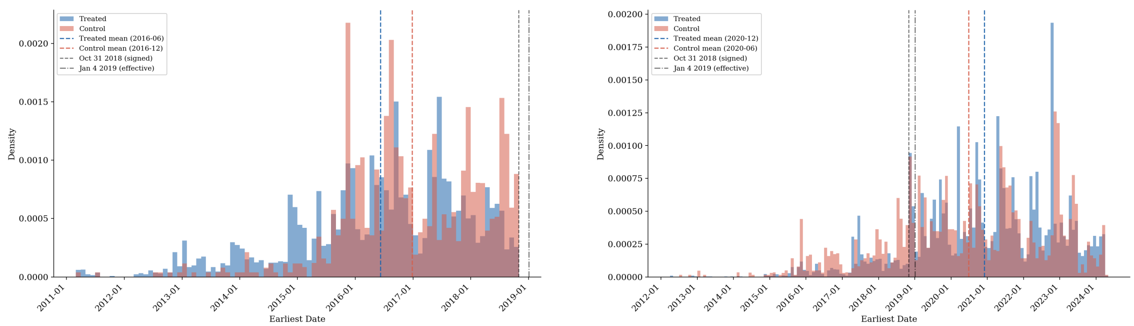

(a) Pre-agreement (May–October 2018) (b) Post-agreement (January–September 2019)

Note: Each panel plots a histogram of the earliest observed listing date for Apple tablet ASINs active in the respective period. *Treated* countries are Germany, Spain, France, Italy, Japan, and the United States; *Control* countries are Canada and Mexico. The rightward shift in the treated distribution between panels indicates that the post-agreement product set in treated markets is concentrated among more recently introduced listings.

Before the agreement, the control group's listings were on average more recently introduced, with a mean introduction date of around December 2016, compared with June 2016 for the treated group. After the agreement, the treated distribution shifts markedly toward newer products through two channels: products introduced before 2017 largely disappeared from treated markets, and a large number of newly introduced models entered the platform. As a result, the average introduction date of the treated group's Apple tablet pool advanced to approximately December 2020, compared with June 2020 in the control group, reversing the pre-agreement gap.

5.2 Buy Box Algorithm and Buy Box Price

Given the reduction in seller participation and the shift in product composition documented above, a natural follow up questions would then be how the agreement affected buyer purchasing. We focus on the Buy Box, which is the default purchase path on Amazon and accounts for the majority of transactions on the platform.[25] Given that, changes in the Buy Box price translate directly into changes in the effective price majority consumers pay, and therefore into consumer welfare. The Buy Box price also reflects the competitive pressure among sellers: when multiple resellers compete aggressively, the Buy Box price tends to align closely with the lowest available offer, whereas a reduction in competition enables the Buy Box to settle at higher levels.

We first examine the average monthly share of Buy Box wins attributed to Amazon for Apple tablets in treated and control markets. As shown in Figure 4a, before the agreement, Amazon's Buy Box share was near zero in both groups, averaging less than 5% across all countries. After

[25] We report the corresponding results for the minimum product price in the Appendix.

the agreement, it spiked to over 40%. This surge is specific to Apple: as shown in Figure 4c, no comparable increase in Amazon's Buy Box share is observed for Samsung tablets over the same period. Moreover, the pattern is consistent across all treated markets, though heterogeneous in magnitude. As shown in Figure 4e, Amazon's share rose from near zero before the agreement, stabilizing at 50–60% in European markets, around 20% in the United States, and near 100% in Japan, while the control group remained near zero throughout.

Figure 4: Impact of the Apple–Amazon agreement on the Amazon share as Buy Box seller of Apple tablets

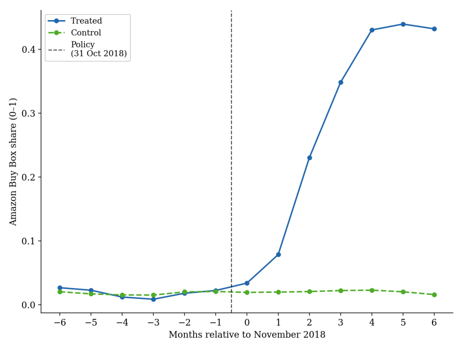

(a) Raw averages: treated vs control

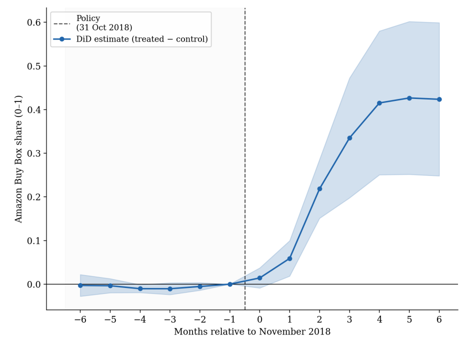

(b) Event Study: Amazon as Buy Box Seller

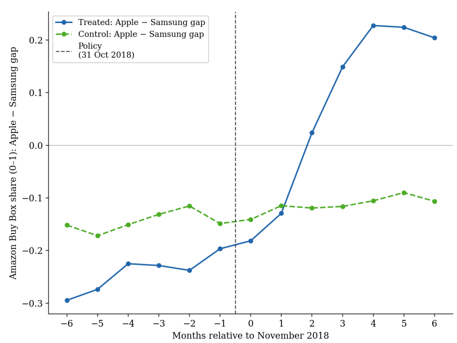

(c) Diagnostic: Apple–Samsung difference

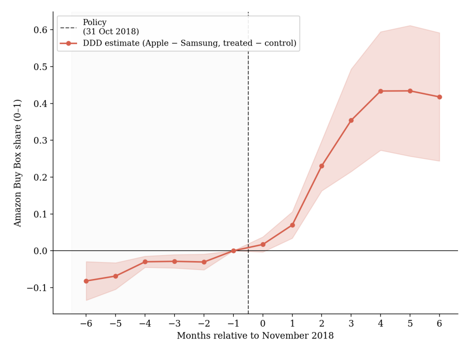

(d) Dynamic Triple-Difference: Apple vs Samsung

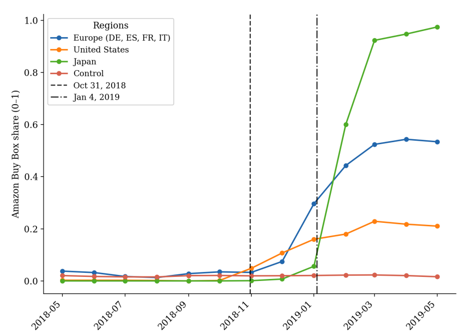

(e) Average monthly Amazon share as Buy Box seller of Apple tablets by region

Note: Figure 4a reports raw averages; Figure 4b shows the event-study estimates; Figure 4c shows the raw diagnostic Apple–Samsung difference; Figure 4d presents the dynamic triple-difference estimates contrasting Apple with Samsung; Figure 4e shows the shares by region.

In Figures 4b and 4d, we report the corresponding event-study and dynamic DDD estimates, which implies that Amazon's share in buy box winner increase by approximately **50%** after the agreement.

Since the Buy Box algorithm is understood to weight fulfillment logistics alongside price, this reallocation of demand toward Amazon need not have been neutral across the third-party sellers it displaced. We therefore decompose the remaining, non-Amazon Buy Box wins into Fulfilled by Amazon (FBA), where the third-party seller ships through Amazon's own warehousing network, and Fulfilled by Merchant (FBM), where the seller ships independently, and re-estimate the event-study and dynamic DDD specifications on each share.

Figure 5: Impact of the Apple–Amazon agreement on the fulfillment composition of Apple tablets' Buy Box winners

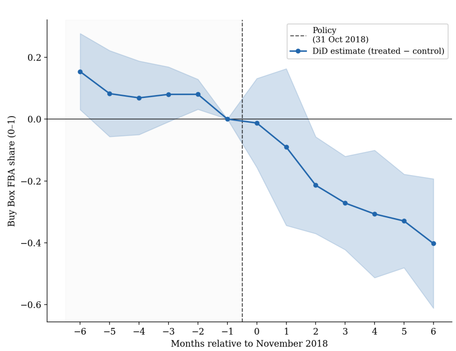

(a) Event Study: FBA share of Buy Box wins

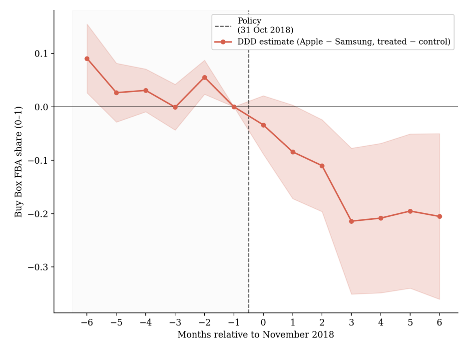

(b) Dynamic Triple-Difference: FBA share, Apple vs Samsung

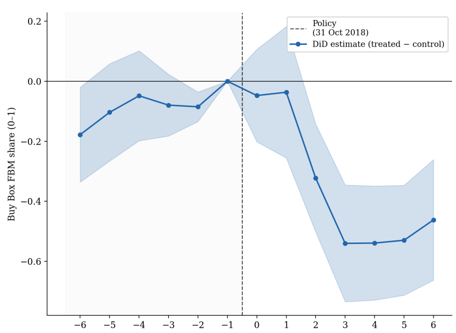

(c) Event Study: FBM share of Buy Box wins

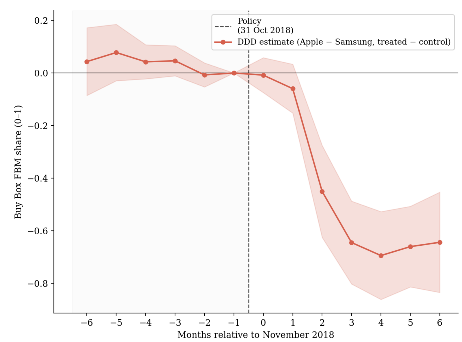

(d) Dynamic Triple-Difference: FBM share, Apple vs Samsung

Note: Figure 5a shows the event-study estimates for the Fulfilled-by-Amazon (FBA) share of Apple tablet Buy Box wins (Apple only, treated vs. control); Figure 5b presents the corresponding dynamic triple-difference estimates contrasting Apple with Samsung; Figures 5c and 5d report the analogous estimates for the Fulfilled-by-Merchant (FBM) share.

As shown in Figures 5b and 5d, the two channels move very differently: the FBM share falls sharply and persistently, by roughly 65–70 percentage points relative to Samsung, with a flat and insignificant pre-trend, while the FBA share declines much more modestly, by about 20 percentage points. This asymmetry suggests that Amazon's Buy Box takeover crowded out merchant-fulfilled sellers far more than Amazon-fulfilled ones.

Besides product diversity, Buy Box price is another key outcome on Amazon marketplaces. It is the price at which the vast majority of Amazon transactions are fulfilled. As such, changes in the Buy Box price translate directly into changes in the effective price paid by most consumers and, consequently, into consumer welfare. The Buy Box price also reflects competitive pressure among sellers: when multiple resellers compete aggressively, the Buy Box price tends to align closely with the lowest available offer, whereas a reduction in competition allows the Buy Box to settle at higher levels.

With the substantial changes in the Buy Box algorithm and seller participation, we find that the average Buy Box price for Apple tablets increased significantly after the agreement. Figures 6b and 6d report the event-study estimates and the dynamic DDD estimates. Both estimations suggest consistent results, showing an Apple-specific increase of approximately 16% increase in the Buy Box price in the months immediately following the agreement.

Figure 6: Buy Box Price Analysis

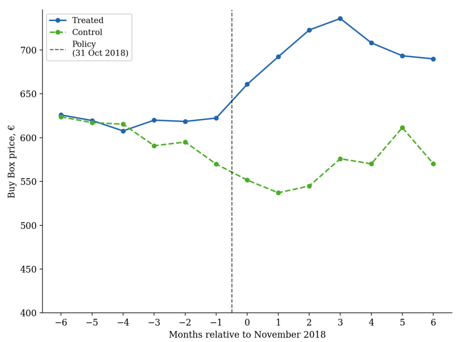

(a) Raw averages: treated vs control

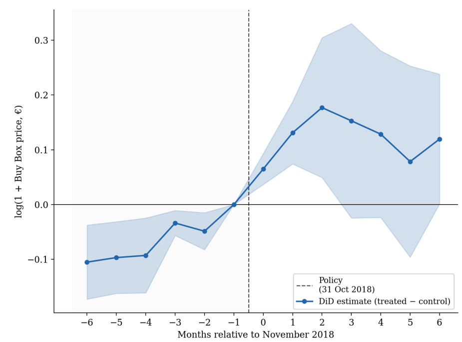

(b) Event Study: Buy Box Price

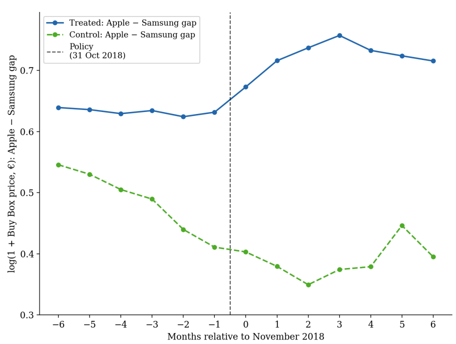

(c) Diagnostic: Apple–Samsung difference

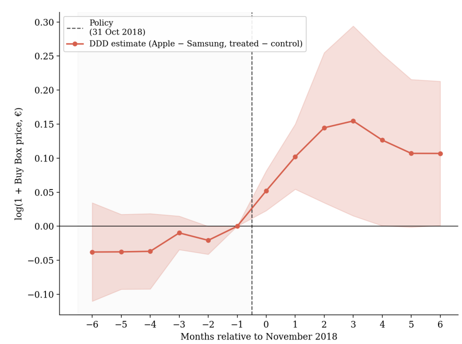

(d) Dynamic Triple-Difference: Apple vs Samsung

Note: Figure 6a reports raw averages; Figure 6b shows the event-study estimates; Figure 6c shows the raw diagnostic Apple–Samsung difference; Figure 6d presents the dynamic triple-difference estimates contrasting Apple with Samsung.

This price increase is not surprising and can arise through several channels. First, the removal of older, lower-priced products shifts the composition of active listings toward newer, higher-priced models. Second, the introduction of new products at higher price points further elevates the average Buy Box price. Third, fewer active resellers reduce intra-brand price competition, weakening the downward pressure that rival sellers had previously exerted. Last, Amazon's expanded role as the primary Buy Box winner means that its pricing—unconstrained by competitive bids from third-party merchants—now determines the effective transaction price for most consumers. In

Section 6, we disentangle these channels.

5.3 Product Quality

We lastly turn to consumer ratings and reviews, which provide a complementary perspective on the effects of the Apple–Amazon agreement. Reviews matter because they not only reflect consumer satisfaction but also influence future purchasing decisions, making them an important dimension of welfare. If the agreement altered the composition of products available on Amazon, changes in average ratings could reveal whether the products that remained in the market were systematically valued more or less by consumers.

Figure 7: Impact of the Apple–Amazon agreement on Apple tablet product ratings

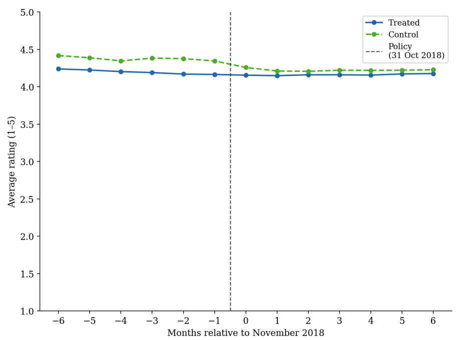

(a) Raw averages: treated vs control

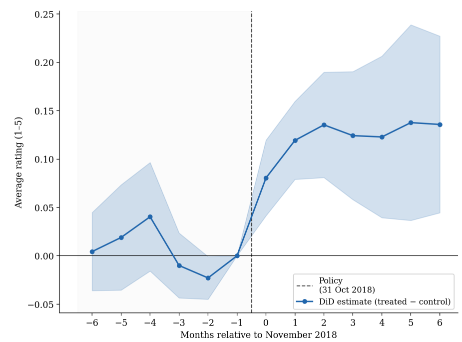

(b) Event Study: product ratings

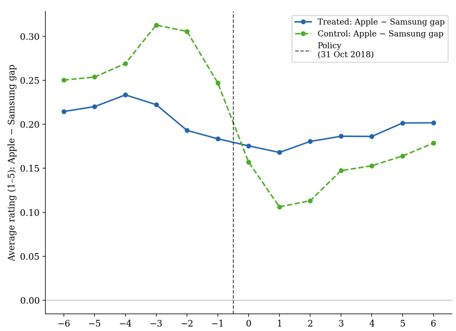

(c) Diagnostic: Apple–Samsung difference

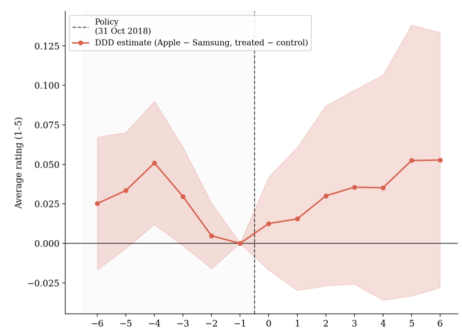

(d) Dynamic Triple-Difference: Apple vs Samsung

Note: Figure 7a reports raw averages; Figure 7b shows the event-study estimates; Figure 7c shows the raw diagnostic Apple–Samsung difference; Figure 7d presents the dynamic triple-difference estimates contrasting Apple with Samsung.

The main findings are summarized in Figure 7. According to the raw averages in Figure 7a, product ratings are high and relatively stable in both treated and control markets throughout the sample period. The event-study estimates in Figure 7b indicate a small increase in treated-market iPad ratings after the agreement. relative to the control, which is driven by the slight decrease in control markets. However, this treated–control comparison does not separate an Apple-specific agreement effect from broader tablet-market changes or brand-level trends. Once we benchmark Apple tablets against Samsung tablets in the DDD specification, the evidence for a quality improvement becomes much weaker. As shown in Figure 7d, the post-agreement DDD coefficients are close to zero, economically small, and statistically imprecise. The diagnostic plot shows no sharp Apple–Samsung break in treated markets relative to the control group.Thus, while average ratings may rise modestly in a simple event-study specification, the more conservative DDD evidence does not support a robust Apple-specific improvement in product quality.

Taken together, the results in this section show that the Apple–Amazon agreement produced a significant change in the composition of the iPad marketplace on Amazon. First, a large share of independent sellers exited the market after the agreement took effect. Second, a substantial portion of older Apple tablets was removed from the platform, sharply reducing the availability of legacy products. Third, the simultaneous introduction of newer products changed the set of options consumers faced, shifting the consideration set toward more recent and generally higher-priced models. Fourth, Amazon became the dominant Buy Box seller for Apple tablets, meaning that its offers were increasingly the default recommendation shown to buyers. Finally, consumers paid higher Buy Box prices on average after the agreement. These patterns indicate that the agreement did not merely change seller identities; it reshaped product availability, recommendation allocation, and the effective prices faced by consumers.

5.4 Robustness Check

We admit that the treated and control markets differ in their baseline level of seller activity. The treated group—the United States, Germany, France, Italy, Spain, and Japan—consists of the world's largest iPad markets, while the control group—Canada and Mexico—represents substantially smaller economies. A level gap in seller counts is therefore expected and largely mechanical.

However, the parallel-trends assumption underlying our DiD estimator requires that treated and control markets share common pre-period *trends*, not common *levels*. A level gap between large and small markets is consistent with valid DiD identification, provided that, absent the agreement, the two groups would have evolved along parallel paths. We verify this directly: the event-study coefficients for seller counts—and for the other outcome variables—are statistically indistinguishable from zero across the pre-agreement periods $k = -6, \ldots, -2$ (see Section 5), providing direct evidence of parallel pre-trends. Additionally, we express all continuous outcome variables in logarithms. Working in log units converts the level difference into a proportional one, which is absorbed by country fixed effects in the panel regression. A market with ten sellers and one with two differ by one unit in log scale regardless of the absolute gap, substantially mitigating the concern that raw numerical disparity drives the post-period estimates.

Moreover, the dynamic triple-difference (DDD) estimator—our preferred specification—largely neutralises concerns about structural differences between treated and control markets. By benchmarking Apple tablets against Samsung tablets within every country and month, the DDD absorbs time-invariant country-level characteristics, including market size and marketplace depth, and identifies the policy effect from the Apple-specific, treated-market break relative to an unaffected brand. The DDD estimates can therefore be interpreted as a within-country, cross-brand difference-in-differences, further insulating the results from market-size confounds.

Lastly, we conduct an in-time placebo test to assess whether the specific timing of the November 2018 agreement drives our estimates. We assign a false policy date of November 2017—exactly one year before the true agreement—and re-estimate the dynamic DDD specification over the window

May 2017 to May 2018, a period that lies entirely before the actual agreement took effect. The identifying logic is straightforward: if structural differences between treated and control markets generated spurious post-period effects at any date, we should observe significant coefficients around the placebo anchor as well; if the November 2018 agreement is the true source of variation, the placebo coefficients should be indistinguishable from zero. Figure 8 reports the results for our four main outcome variables. In every panel, the DDD coefficients at $k = 0, 1, \ldots, 6$ relative to November 2017 are close to zero and statistically insignificant, and the pre-period trends are flat. This null result strongly corroborates the view that the large post-period effects documented in Section 5 reflect the specific timing of the Apple–Amazon agreement rather than pre-existing structural divergence between treated and control markets.

Figure 8: In-Time Placebo: Dynamic DDD Estimates with Fake Policy Date (November 2017)

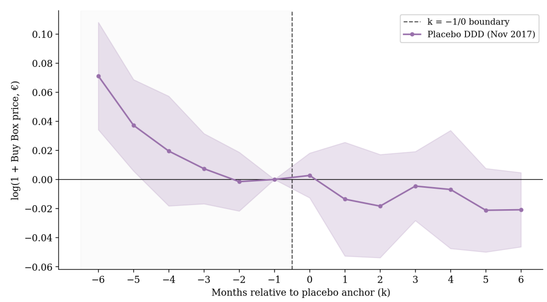

(a) Buy Box Price

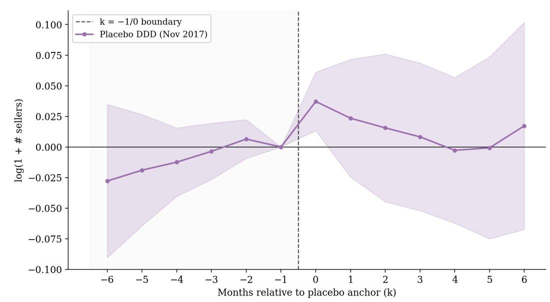

(b) Number of Sellers

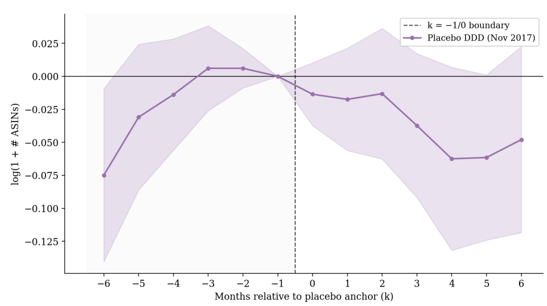

(c) Number of ASINs

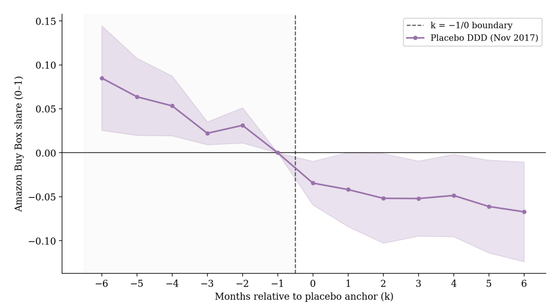

(d) Amazon Buy Box Share

Note: Each panel reports dynamic triple-difference (DDD) estimates assigning a placebo policy date of November 2017. The estimation window is May 2017–May 2018, a period lying entirely before the actual Apple–Amazon agreement. $k = 0$ denotes the placebo anchor month; $k = -1$ (October 2017) is the omitted baseline. The dashed vertical line marks the $k = -1/k = 0$ boundary. Shaded bands are 95% confidence intervals. Standard errors are clustered by country.

6 Underlying Mechanism

The results in the previous section establish that the Apple–Amazon agreement led to significant changes in the structure of the iPad marketplace: fewer independent sellers, a contraction in the number of available products, and higher average Buy Box prices. This section investigates the mechanisms behind these changes by examining which products exited, which remained, and which entered after the agreement, as well as what distinguished these groups.

6.1 Impact of Seller and Product Exit and Product Introduction

To understand how the agreement reshaped product availability, we classify Apple tablet ASINs into three groups based on their listing histories. *Exiters* are products that were first observed before October 2018 and exited the marketplace between November 1, 2018 and January 4, 2019—the window during which the brand-gating provisions of the agreement took effect. *Stayers* are products that were first observed before October 31, 2018 and remained available after January 4, 2019. *Entrants (New ASINs)* are ASINs first observed after November 1, 2019, representing products introduced following the agreement.

Apple and Amazon justified the agreement publicly on quality grounds—the stated rationale was to protect consumers from counterfeit goods and ensure consistent product quality by restricting sales to authorized resellers. We first ask whether the data support this claim: did the products and sellers that exited have worse quality indicators than those that remained? We then ask a complementary question: if not quality, what actually distinguished the exiters?

6.1.1 Product Exit

Figure 9: Product Rating Analysis

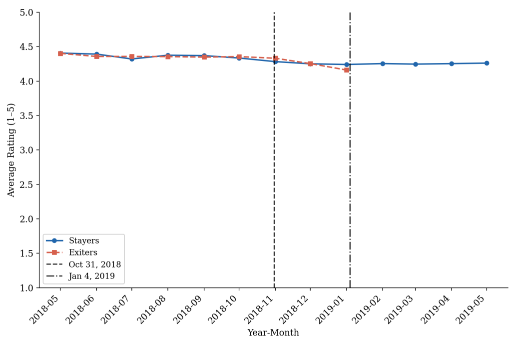

(a) Control countries: exiters vs stayers

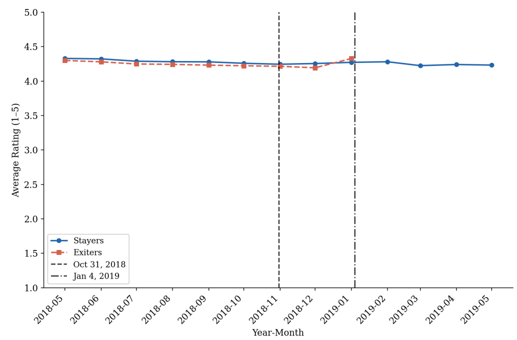

(b) Treated countries: exiters vs stayers

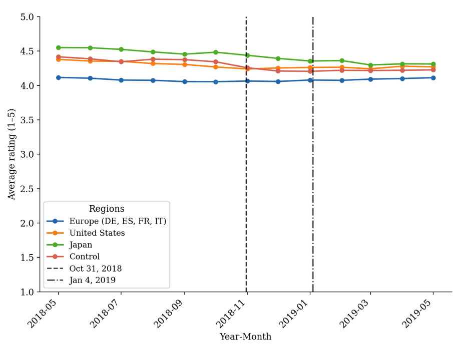

(c) Regional averages

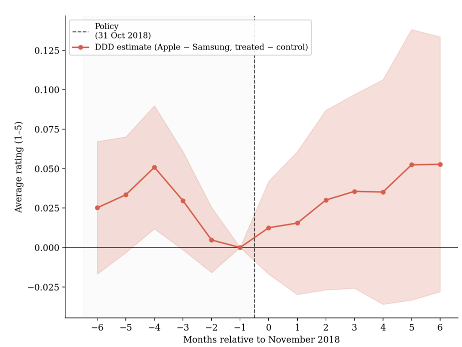

(d) Dynamic triple-difference estimates

Note: Figure 9a and Figure 9b plot monthly average ratings for exiters and stayers in control and treated markets, respectively. Figure 9c reports average ratings across regions. Figure 9d shows dynamic triple-difference estimates comparing Apple and Samsung tablets in treated vs. control countries.

Quality screening: product ratings. Figures 9a and 9b plot the average monthly ratings of Apple tablets separately for exiters and stayers. In the control countries, exiters and stayers displayed nearly identical ratings before October 2018, with averages clustered around 4.3. After that point, exiters' ratings declined somewhat, but this pattern is best interpreted as mechanical: once the exit cohort became small, average ratings fluctuated more due to the diminished sample size. A similar pattern appears in the treated countries. Prior to October 2018, exiters and stayers followed almost indistinguishable trajectories, averaging around 4.2–4.3; after the agreement, exiters' ratings exhibited more volatility, but these changes largely reflected noise rather than systematic differences. The ratings of exiters and stayers were not statistically significantly different in either group of countries, suggesting that the agreement did not selectively remove lower- or higher-rated products.

Across all markets, Figure 9c shows that ratings remained remarkably stable, with no break around the signing or implementation of the agreement. In Europe and the United States, average ratings fluctuated within a narrow band of 4.0–4.2, while Japan maintained persistently higher ratings of about 4.5–4.6. The dynamic triple-difference estimates in Figure 9d confirm this picture: pre-period coefficients are centered around zero, and post-agreement effects are economically negligible—on the order of 0.05–0.1 points on a five-point scale—with wide confidence intervals consistent with no meaningful change.

As a further robustness check, we also examined the text of product reviews using a multilingual transformer model that predicts star ratings from review content. Restricting the sample to the pre-agreement window (January–October 2018) to avoid contamination, we find that exiters and stayers had nearly identical sentiment profiles: exiters registered an average predicted score of 4.32 and an actual rating of 4.49, compared with 4.25 and 4.44 for stayers. Table 2 presents the results. Statistical tests (Welch t-test, $p = 0.12$; Mann-Whitney U, $p = 0.06$) show no significant difference between the two groups. These findings provide little support for Apple and Amazon's quality-based justification: brand gating did not selectively remove lower-quality products from the marketplace.

Table 2: Sentiment analysis

Group	Reviews	Avg pred. stars	Avg conf.	Positive	Neutral	Negative	Avg true
Exiters	1,219	4.321	0.695	83.3%	4.1%	12.6%	4.488
Stayers	3,215	4.253	0.689	81.7%	4.4%	13.8%	4.435

Notes: (1) *Exiters* are products that exited the platform between October 31, 2018 and January 4, 2019; *Stayers* remained listed. (2) *Avg pred. stars* is the model-predicted star score on a 1–5 scale derived from text sentiment; *Avg true* is the observed star rating on the platform. (3) *Avg conf.* is the average model confidence (probability of the assigned class) across reviews. (4) *Positive/Neutral/Negative* are the shares of reviews classified into each sentiment category; percentages may not sum to 100% due to rounding.

Cohort selection: which products exited? Since the evidence rules out quality-based screening, we ask what did in fact distinguish exiters from stayers. We examine product entry cohorts, measured by the earliest observed date of each ASIN in our data.

Figure 10: Distribution of Entry Cohorts: Exiters vs. Stayers

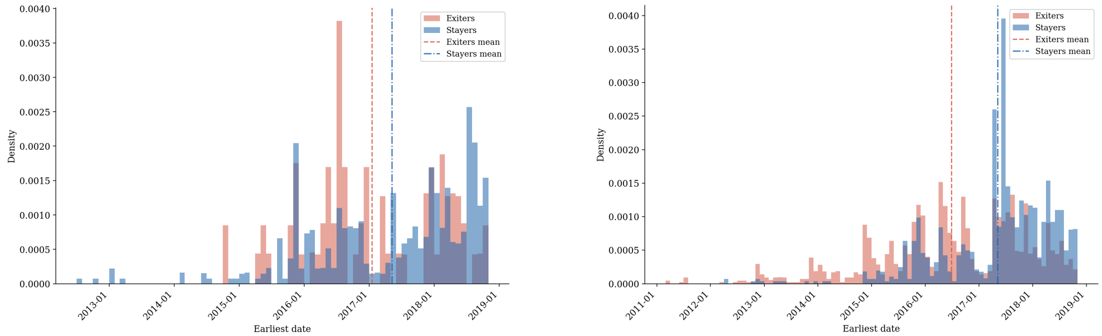

(a) Control countries (b) Treated countries

Note: The figure plots kernel density estimates of earliest observed listing dates for Apple tablet ASINs. Exiters are defined as products that exited the marketplace during the implementation window of the agreement. Stayers are products that remained listed after implementation. Vertical lines indicate mean entry dates for each group.

Figure 10 plots kernel density estimates of entry dates separately for exiters and stayers in treated and control markets, allowing us to compare not only mean entry timing but the full distribution of cohorts. In both treated and control countries, exiters entered earlier on average than stayers, consistent with standard product life-cycle dynamics: older products are generally more likely to exit. However, the separation between the two distributions is substantially stronger in treated markets.

In treated countries, the distribution of exiters is clearly shifted to the left relative to stayers, with greater mass in early entry cohorts and a thicker left tail. Stayers are disproportionately concentrated among more recent entrants, particularly in the years immediately preceding the agreement. Importantly, the divergence extends across much of the support of the distribution, not merely at the mean. By contrast, in control markets the distributions overlap considerably more, and tail differences are modest.

These patterns indicate that the agreement induced cohort-based selection rather than quality-based selection. In treated markets, products drawn from earlier entry cohorts were disproportionately removed, while more recent entrants were substantially more likely to remain listed. The weaker separation in control countries suggests that this pattern cannot be explained solely by natural aging or global product cycles. Instead, the evidence is consistent with the agreement accelerating the exit of earlier-generation products in treated markets, consolidating distribution around newer models without improving observable product quality.

Figure 11: Buy Box Price: Exiters vs Stayers

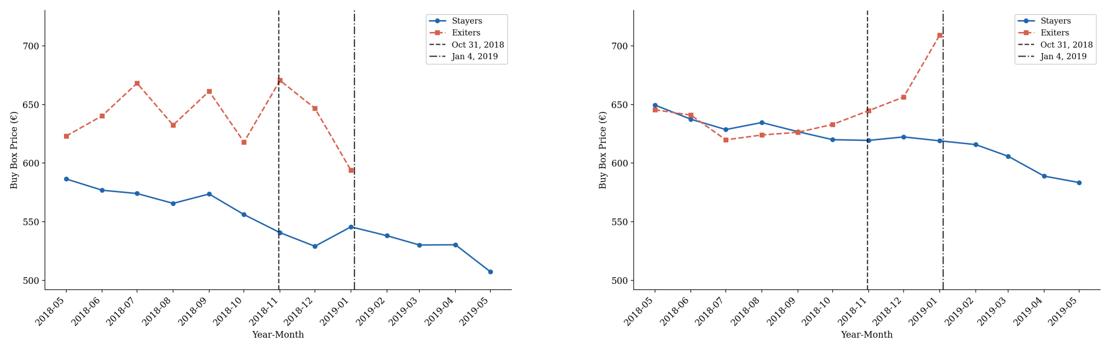

(a) Control countries

(b) Treated countries

Note: Figure 11a and Figure 11b plot the monthly average Buy Box prices for Apple tablets in control and treated marketplaces, respectively.

Composition effects on prices. Figure 11 examines whether the products that exited were systematically different in their pricing from those that survived. If exiters tended to be lower-priced, their disappearance would have amplified the observed price increases by removing the more competitive fringe of the market.

In the control group, exiters were systematically priced above stayers throughout the sample period—a pattern consistent with standard market dynamics in which very high-priced products face limited demand and are more likely to disappear.

The data patterns look different in the treated group. As in the controls, exiters in the treated group were also priced similarly to the stayers almost till the adgreement implementation. Exiters in treated markets exhibited a sharp spike in average Buy Box prices in the months surrounding the agreement's implementation. The sellers most likely to exit were those competing aggressively at lower price points, so their removal left behind higher-priced offers and introduced greater fluctuations in the observed Buy Box price. The price volatility of the exiters is thus consistent with the mechanism: cheaper resellers were forced off the platform, the remaining offers temporarily set higher prices, leading to sudden jumps in the Buy Box price.

Two dimensions of the agreement's impact on prices stand out. First, by reducing the presence of lower-priced resellers, the agreement altered the composition of available products and narrowed the difference in price levels between exiters and stayers. Second, the timing of exits in treated markets coincided with spikes in Buy Box prices, underscoring how the sudden withdrawal of cheaper offers translated directly into higher effective prices faced by consumers.

6.1.2 Seller Exit

Quality screening: seller ratings. Seller ratings provide a complementary perspective on whether the agreement improved the quality of the marketplace by filtering out weaker sellers. In the control markets (Figure 12a), sellers who exited had consistently lower ratings than those who remained—exiters started with an average rating of approximately 90, while stayers averaged 92–94—suggesting that natural seller exit in the control group was correlated with lower performance.

Figure 12: Seller Rating Analysis

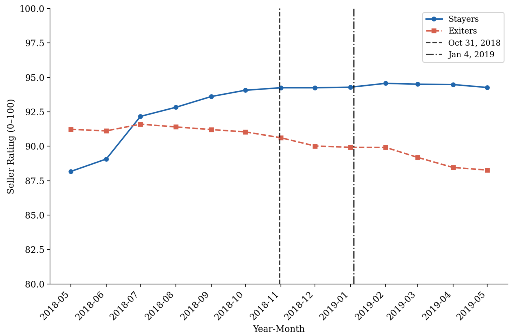

(a) Control countries: exiters vs stayers

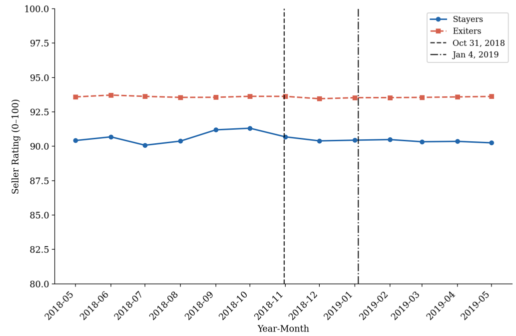

(b) Treated countries: exiters vs stayers

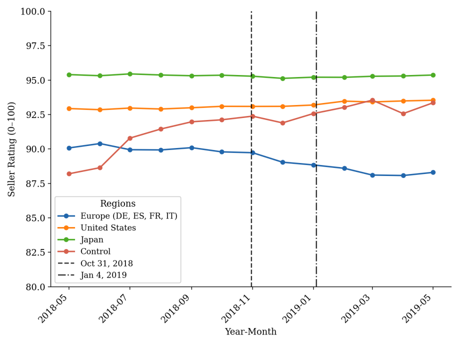

(c) Ratings across regions

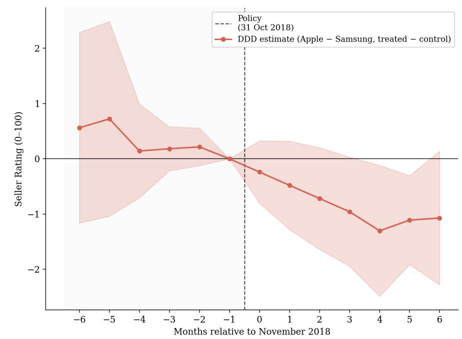

(d) Dynamic triple-difference estimates

Note: Figures 12a and 12b compare ratings for exiters and stayers in control and treated markets, respectively. Figure 12c reports average seller ratings by region over time. Figure 12d shows dynamic triple-difference estimates comparing seller ratings in treated vs. control countries.

The treated group tells a strikingly different story. Figure 12b shows that, in treated markets, exiters did not exhibit systematically lower ratings than stayers; indeed, exiters maintained even higher ratings than stayers throughout the sample period. This reversal implies that the sellers removed by the brand-gating agreement were not the lower-rated ones. Figure 12c reinforces this: in control countries, average seller ratings improved steadily over the sample period, consistent with the gradual exit of weaker sellers, while in treated countries ratings remained flat or declined slightly, indicating no improvement in seller quality following the agreement. The dynamic triple-difference estimates in Figure 12d confirm the absence of any significant relative improvement in seller ratings for Apple tablets in treated versus control markets.

Together with the product ratings evidence, the seller rating analysis underscores that brand gating did not operate as a quality-screening mechanism. The sellers expelled from treated markets were not lower-rated than those who remained—if anything, the opposite holds.

Which sellers exited? Since the evidence rules out quality-based seller screening as well, we ask what did in fact distinguish exiters from stayers. The sellers who exited the iPad market following the Amazon–Apple agreement were not simply fringe operators. Along several dimensions they were comparable to or even stronger than those who remained. Apart from the higher quality, FBA adoption was nearly identical: 50% of exiters used Fulfilled by Amazon vs 53% of stayers, ruling out fulfillment method as a main differentiating factor. The two groups were also similarly committed to the iPad category: exiters were active for 11 months on iPads pre-policy vs 11 months for stayers – a gap of barely one month. Their iPad share of total storefront was likewise close: 26% for exiters vs 23% for stayers, meaning iPads were almost equally central to both groups' business.

The sharpest contrast is in seller scale. Stayers were much larger merchants overall,they had an average of 28,294 total storefront ASINs vs only 4,827 for exiters – nearly a 6 times difference. This is also visible in iPad catalog breadth: stayers listed an average of 38 distinct iPad ASINs vs 17 for exiters. The CDFs confirm this (see Figure 13 below): the stayer distribution has more mass in the upper tail in terms of number of storefront ASINs and iPad ASINs. Exiters were more niche sellers — highly rated, iPad-focused, but operating at smaller scale.

Figure 13: Treated countries: Exiters vs. Stayers among Sellers

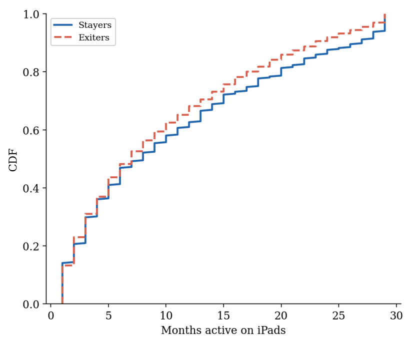

(a) Months active on iPad

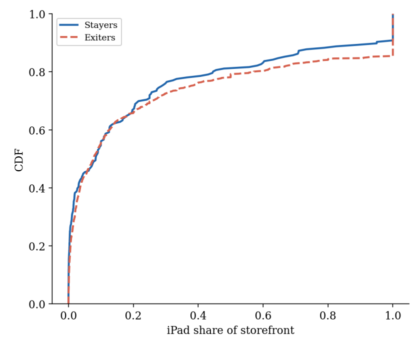

(b) iPad share of Storefront

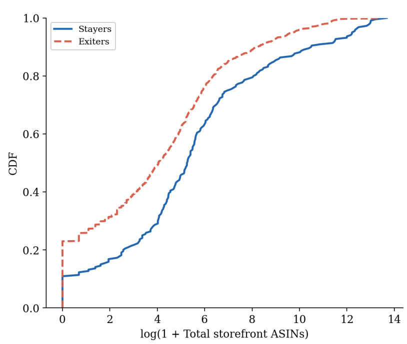

(c) Total Storefront ASINs

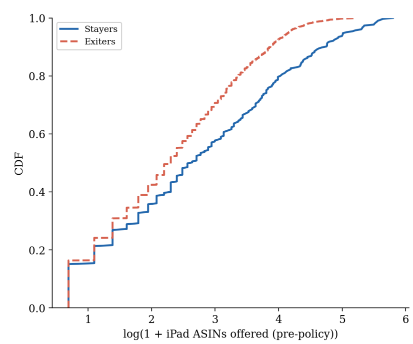

(d) Total iPad ASINs

Note: Figure 13a and Figure 13b show that exiters and stayer in treated markets were similar in terms of the number of mmonths active in iPad market as well as iPad specialization. Figure 13c shows that stayers are bigger and Figure 13b shows that stayers also offered more iPads.

The buy box winning shares provides extra evidence on the importance of exiting sellers and competitive effects. Figure 14 shows that before the policy, exiters collectively help approximately 55–60% of iPad buy box share in treated countries compared to 25–30% for the stayers. After the policy takes effect (post January 4, 2019), the picture reverses sharply: exiters' share collapses toward zero, while stayers absorb nearly 90% of buy box. New entrants ("Other") capture only a modest residual share (around 10%), confirming that the buy box was not redistributed to fresh competitors but concentrated among Amazon and the pre-policy incumbents who qualified under the new brand-gating rules.

Figure 14: Buy Box Shares: Exiters vs. Stayers

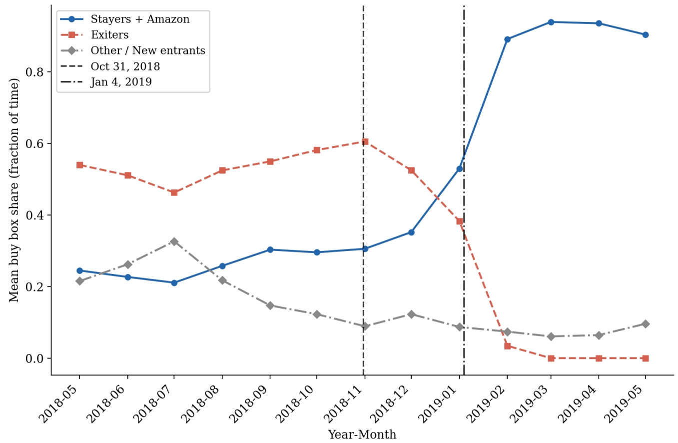

Note: The figure compares stayers and exiters in treated marketplaces in terms of teh buy box shares. The dashed and dash-dotted vertical lines indicate October 31, 2018 and January 4, 2019, respectively.

The pattern is consistent and economically meaningful: Amazon's agreement with Apple did not selectively remove low-quality or marginal sellers. Those who exited were better-rated than those who stayed, equally reliant on FBA, and similarly committed to the iPad category — they were simply smaller. The policy effectively installed a scale threshold that small high-quality specialists could not clear, concentrating buy box ownership among large incumbents and Amazon itself. This is difficult to reconcile with a quality-enhancing justification for the agreement.

6.1.3 Product Introduction

The preceding analysis showed that the agreement was followed by a wave of exits concentrated among older Apple tablet ASINs. We now examine the complementary side of the market: the products that entered *after* the agreement was implemented. Entrants are defined as ASINs first observed on Amazon after November 1, 2018—products whose earliest listing date falls after the start of the exit wave. Figure 15 characterizes these new entrants by comparing their Buy Box prices and product ratings to those of stayers throughout the sample period.

Figure 15: New Products Entry

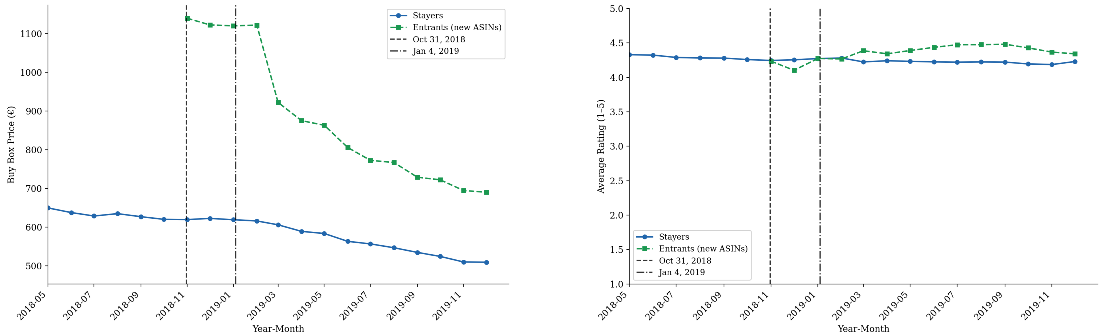

(a) Buy Box price: stayers vs entrants

(b) Rating: stayers vs entrants

Note: Figure 15a and Figure 15b plot monthly average buy box prices and ratings for stayers and entrants treated markets, respectively.

Panel (a) reveals a striking price gap between entrants and stayers at the time of entry. New ASINs first appeared in treated markets at average Buy Box prices of approximately 1,140 euros —nearly 80 percent above the stayer average of around 620 euros at the same point in time. This premium reflects the composition of the entrant cohort: the products being introduced after the agreement was implemented were primarily Apple's latest models, which command launch prices that substantially exceed those of older iPad lines still active as stayers. Entrant prices declined over the course of 2019 as these products aged, the initial cohort of early adopters was satisfied, and within-listing competition intensified. By the end of 2019, the average entrant Buy Box price had fallen to approximately 690 euros, while stayer prices had also drifted downward to around 510 euros. Even after this convergence, however, entrants remained priced noticeably above stayers throughout the entire observation window.

Panel (b) shows that this price premium is not driven by substantial quality differences. Both entrants and stayers maintained ratings in the 4.2–4.5 range, with entrants achieving ratings that were slightly above, those of stayers from mid-2019 onward. The higher prices of entrants reflect their position in the product cycle: newly released models command a premium because they represent the current technology frontier, not because they are significantly higher ranked relative to the models that were removed.

A complementary way to read the seller-exit results above is to ask not only how many sellers remain on a listing, but how differently the surviving sellers price it. If a product was, before the agreement, served by resellers sourcing through different channels, such as authorized wholesale, parallel import, or informal gray-market supply, their prices would be expected to diverge, and the degree of that divergence is itself informative about how many meaningfully distinct competitive channels were active. A narrowing of price dispersion for the same product is therefore a complementary measure of eroding competitiveness that does not rely solely on counting sellers.

We measure this dispersion by pooling, for each ASIN and country, all seller-day offers observed within a calendar month and computing the coefficient of variation of price across that pooled set.[26]

[26] Offers are restricted to new-condition listings and converted to euros using the same exchange-rate and inflation adjustment applied throughout the paper.

We use the coefficient of variation, the standard deviation scaled by the mean, rather than the raw variance, for two reasons. First, it is scale-free: raw variance, measured in squared euros, mechanically scales with a product's price level, so a higher-priced iPad model would appear more "dispersed" than a cheaper one even if sellers disagreed on price by exactly the same proportion, an artifact the coefficient of variation removes, making dispersion comparable across products, cohorts, and time. Second, pooling an entire month's offers into a single calculation, rather than computing dispersion separately for each day and averaging across days, avoids letting one day's idiosyncratic price gap dominate the month, a real risk once seller counts fall to two or three following the agreement. Figure 16 applies this measure separately to stayers and entrants in treated markets, starting May 2018.[27]

Figure 16: Offer Price Dispersion: Stayers vs Entrants (Treated Markets)

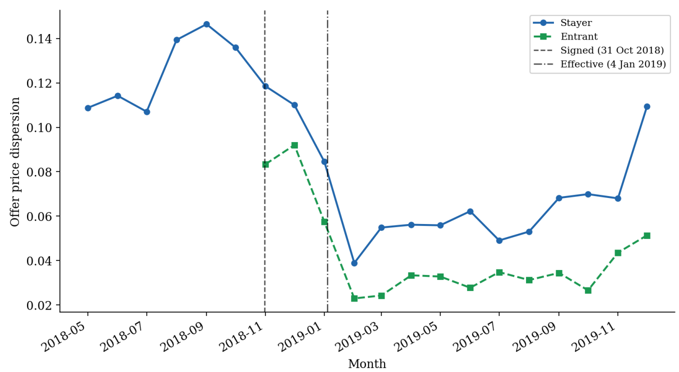

Note: Monthly coefficient of variation of price across all seller-day offers pooled within each ASIN-country-month, averaged across ASINs, for stayers and entrants in treated markets. Dashed and dash-dotted vertical lines indicate October 31, 2018 and January 4, 2019.

The pattern traces a clear ordering. Before the agreement, stayers show substantial dispersion, with a coefficient of variation of roughly 11–15 percent from May through December 2018, consistent with resellers sourcing through different channels pricing the same product quite differently. Dispersion remains elevated through the end of 2018 and falls sharply beginning in February 2019, settling at roughly 4–7 percent for most of the remainder of the sample. Entrants show even less dispersion, settling at roughly 2–4 percent for most of 2019 after an initial launch-period spike in November and December 2018, consistent with products that never operated under the earlier multi-channel supply structure and whose pricing, from the outset, was concentrated among the small number of sellers Amazon permitted. This ordering, pre-agreement stayers most dispersed, post-agreement stayers intermediate, and entrants least dispersed, is consistent with a broad narrowing of the scope for independent price-setting following the agreement. Because this measure pools Amazon-entered and never-entered products together, it captures a market-wide average rather than isolating Amazon's role; the more granular comparison documented later in this section shows that price stabilization is in fact concentrated among the subset of stayers Amazon subsequently entered, while never-entered stayers continue to see prices decline. Amazon's Buy Box capture is therefore likely an important contributor to the dispersion decline shown here, but not the sole one.

[27] Entrants necessarily first appear only after November 2018, so their series begins later than stayers'.

Together, the exit, entry, and dispersion patterns document a coherent compositional shift. The agreement reshuffled the Apple tablet assortment available on Amazon: older, lower-priced models were removed from the platform, while newer, higher-priced models took their place. This reshuffle raised the average price level of the consideration set independently of any change in list prices or retailer margins. The dispersion evidence adds a second dimension to this shift: cross-seller price disagreement on the same product fell sharply after the agreement, and entrants, which never experienced the earlier, less restricted marketplace, show the least price disagreement of all. The mechanism therefore operates through product mix as well as through a narrowing of price variety among listed sellers: by restricting which products are eligible for listing and how much dispersion survives among the sellers who remain, the Apple–Amazon agreement effectively elevated the price floor of the market.

6.2 Role of Amazon

The preceding analysis showed that brand gating narrowed price dispersion for the pooled population of stayers, without identifying whether this narrowing on its own raised the average price level or whether that required Amazon's own participation as Buy Box winner. We now disaggregate by Amazon's entry decision to answer this question directly.

6.2.1 Amazon Strategy

Product Entry. Amazon's participation in the post-agreement marketplace was far from passive: the platform actively selected which Apple tablet products to enter as a seller. To illustrate, we classify products by Amazon's entry status over the full sample window. *Amazon-entered* products are those that Amazon did not offer before October 2018 but began offering after that date, i.e., products for which Amazon's entry coincided with or followed the agreement. *Never-entered* products are those that Amazon did not offer at any point in the sample window.

Figure 17 documents the selection pattern across these two groups. Panel (a) plots the entry date distribution for all ASINs in treated countries, split by whether Amazon subsequently entered as a seller. The distribution for Amazon-entered products is heavily right-tailed, concentrated among ASINs introduced after January 2019. A large share of the newly listed products that appeared on the platform immediately following the agreement's implementation were entered by Amazon, while products introduced before the agreement and not chosen by Amazon cluster in earlier entry cohorts. Panel (b) restricts to stayers, defined as products present both before and after the agreement, and shows the same pattern within this subgroup: even among products that survived the exit wave, Amazon selectively entered those with more recent introduction dates, leaving the older stayers to be sold exclusively by remaining third-party sellers.

Figure 17: Distribution of Entry Dates

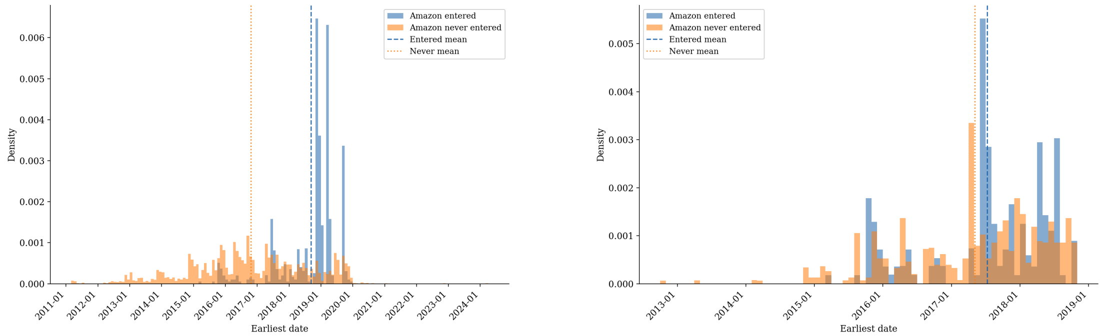

(a) All ASINs in Treated Countries

(b) Stayers in Treated Countries

Note: Figure 17a and Figure 17b plot kernel density estimates of earliest observed listing dates for stayers and newly introduced Apple tablet ASINs in treated and control marketplaces, respectively, split by whether Amazon entered as a seller.

The products Amazon chose to enter were not only newer but also substantially more popular and more expensive. Figure 18 tracks each outcome across four groups. The sample is first split by product cohort: *stayers* are products first observed before October 31, 2018 that remained available after January 4, 2019; *new ASINs* are products first observed after October 31, 2018. Each cohort is then further divided by Amazon's entry status as defined above: whether Amazon began offering the product after October 2018 (*Amazon-entered*) or never offered it within the sample window (*never-entered*). This yields four groups, represented by four lines: the blue solid line (Amazon-entered stayers), the orange solid line (never-entered stayers), the green dashed line (Amazon-entered new ASINs), and the purple dashed line (never-entered new ASINs).

Figure 18a shows seller per ASIN counts over time. Among stayers, Amazon-entered products averaged around 17 active sellers before the agreement, more than three times the approximately 5 sellers per never-entered stayer, indicating that Amazon concentrated its entry on high-demand products that had previously attracted intense third-party competition. After the agreement, brand gating removed unauthorized resellers across the board. Amazon-entered stayers fell sharply from 17 to around 2 sellers per ASIN, while never-entered stayers declined only modestly from 5 to around 3, reflecting their smaller initial seller base. Among newly introduced ASINs, both Amazon-entered and never-entered products begin close to 2–3 sellers and remain flat throughout, consistent with new products entering an already-restricted market.

As for the price, it can illustrated in Figure 18b. Among stayers, Amazon-entered products were already priced around 700 euros before the agreement, compared with approximately 550 euros for never-entered stayers, confirming that Amazon's entry strategy targeted higher-value products. For newly introduced ASINs, Amazon-entered products spike to around 1,250 euros around the time of the agreement, reflecting the introduction of premium models, before stabilizing at approximately 850–950 euros. Never-entered new ASINs also spike briefly to around 990 euros but then collapse to below 500 euros and eventually converge toward the never-entered stayer level, consistent with price competition among third-party sellers on products Amazon did not enter.

Figure 18: Amazon Entry and Market Outcomes: Sellers and Buy Box Prices

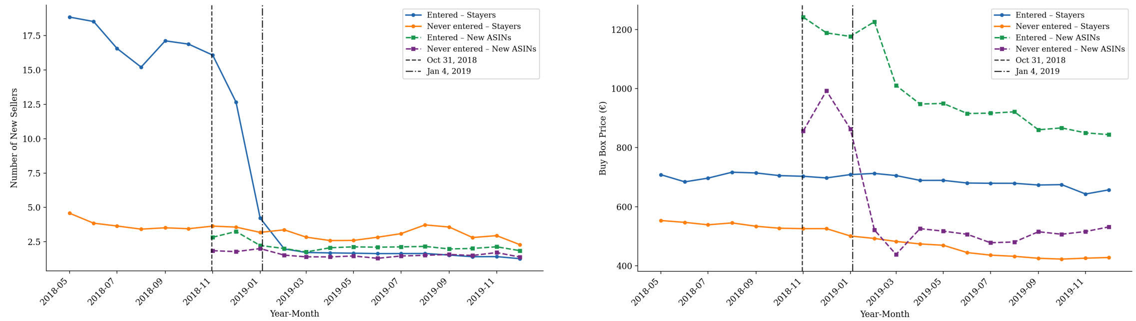

(a) Number of new sellers

(b) Buy Box price

Note: The figure compares stayers and newly introduced Apple tablet ASINs in treated marketplaces, separately by whether Amazon entered as a seller. The dashed and dash-dotted vertical lines indicate October 31, 2018 and January 4, 2019, respectively.

Self-preferencing in Buy Box Algorithm. Once Amazon enters a product, it captures Buy Box assignment at a remarkably high rate. Figure 19 shows Amazon's Buy Box share over time for stayers and newly introduced ASINs. Precisely, for stayers that Amazon entered, its Buy Box share rises to approximately 70% after the agreement. For newly introduced products entered by Amazon, the share approaches 100%.

Figure 19: Amazon Entry and Buy Box Assignment

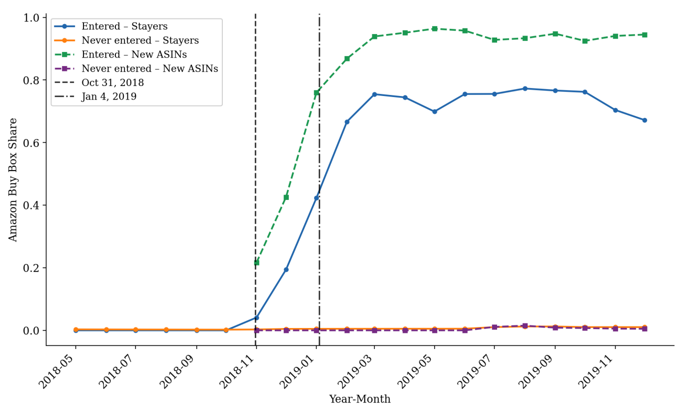

Note: The figure compares stayers and newly introduced Apple tablet ASINs in treated marketplaces, separately by whether Amazon entered as a seller. The dashed and dash-dotted vertical lines indicate October 31, 2018 and January 4, 2019, respectively.

This dominance is not mechanically explained by a lack of competition. As shown in Figure 18a, Amazon-entered stayers retain more third-party sellers than new Amazon-entered products, yet Amazon's Buy Box win rate is *lower* for stayers than for new products. Intuitively, more competing sellers should make it harder for Amazon to win the Buy Box, yet the data show the opposite: Amazon's win rate is actually lower for stayers, where more sellers remain, than for newly introduced ASINs, where competition is thinner. Amazon still commands around 70% of Buy Box assignments even on the more competitive stayer products. This suggests that Amazon's Buy Box dominance is not simply a mechanical outcome of being the only or cheapest seller, but points to self-preferencing in how the algorithm allocates the Buy Box.

Fulfillment parity does not explain the gap either. Section 5 shows that among the remaining third-party sellers, those fulfilling through Amazon's own warehousing network (FBA) retained Buy Box share far better than merchant-fulfilled (FBM) sellers, confirming that fulfillment logistics is a meaningful input into the algorithm. Yet Amazon's own win rate of 70–100% substantially exceeds the retention enjoyed by FBA third-party sellers, who share the same fulfillment channel but not Amazon's ownership of the platform. If fulfillment logistics alone drove Buy Box allocation, Amazon and FBA third-party sellers should fare similarly; the persistent gap between them indicates that Amazon's own listings receive additional algorithmic priority beyond what fulfillment parity would predict.

6.2.2 Impact of Amazon Strategy

To disentangle the impact of Amazon's entry from the mechanical effect of brand gating, we compare outcomes for products with and without Amazon entry, holding the brand gating treatment fixed. The logic is that brand gating applies uniformly to all Apple products in treated markets: every product lost unauthorized third-party sellers. Any difference in outcomes between products that Amazon chose to enter and those it did not must therefore reflect the additional impact of Amazon's own strategic behavior, namely its selective entry and its control over the Buy Box algorithm, rather than the contractual restriction alone. This within-treatment comparison identifies the incremental role of algorithmic demand steering on top of vertical foreclosure.

Buy Box price. Figure 18b reveals a striking divergence in Buy Box prices depending on whether Amazon entered. For stayers that Amazon never entered, prices decline gradually after the agreement, consistent with natural product aging: as older models lose appeal, surviving third-party sellers compete to clear inventory. By contrast, for stayers that Amazon entered and self-preferenced in the Buy Box, prices stabilize rather than fall. Amazon's dominant Buy Box position insulates these products from the downward price pressure that would otherwise accompany an aging product cycle.

The pattern is even sharper for newly introduced products. New ASINs offered by Amazon are priced substantially above existing stayers and maintain those elevated prices over time, as Amazon's Buy Box control shields them from competition. New ASINs not offered by Amazon, however, quickly shed their initial price premium: their Buy Box prices decline rapidly and within months fall below those of Amazon-entered stayers, exposing them to the full force of third-party price competition.

Taken together, these patterns show that seller exit alone does not generate sustained price increases. Even among stayers that lost a large share of their seller base after brand gating, Buy Box prices did not rise substantially. What drove the price changes following the agreement were Amazon's entry and its algorithmic self-preferencing that followed. The aggregate price increase thus operates through two channels: Amazon's self-preferencing stabilizes prices for the stayer products it enters, and its introduction of new products at substantially higher price points, combined with Buy Box control, keeps those prices elevated, collectively raising the average Buy Box price for Apple tablets in treated markets.

Figure 20: Amazon Entry and Product Outcomes: Ratings and Reviews

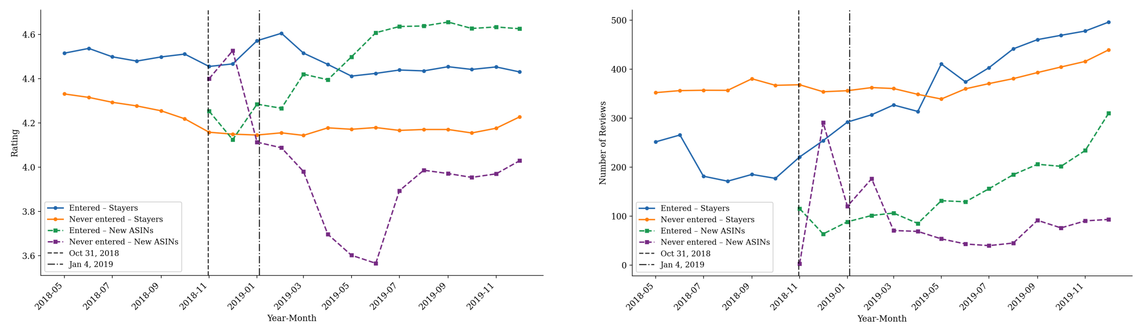

(a) Rating

(b) Number of reviews

Note: The figure compares stayers and newly introduced Apple tablet ASINs in treated marketplaces, separately by whether Amazon entered as a seller. The dashed and dash-dotted vertical lines indicate October 31, 2018 and January 4, 2019, respectively.

Product ratings and reviews. Figure 20 documents a further dimension of Amazon's algorithmic advantage. As shown in Figure 20a, ratings for stayer products remain stable across both groups, regardless of Amazon's entry. Among newly introduced products, however, a sharp divergence emerges. New ASINs offered by Amazon rapidly accumulate strong ratings: Amazon's self-preferencing channels the majority of transactions to its own listings, so these products collect reviews quickly and their average ratings improve over time. New ASINs not offered by Amazon, by contrast, receive ratings that drift downward and exhibit more volatility. This likely reflects the smaller volume of transactions flowing to non-Amazon listings and potential heterogeneity among third-party seller quality.

Figure 20b confirms this divergence in consumer engagement. Amazon-entered products, both stayers and new ASINs, accumulate substantially more reviews over time than comparable products not entered by Amazon, and the gap widens steadily after the agreement. Since review volume is a strong signal of transaction volume, this pattern is consistent with the Buy Box channeling demand almost exclusively toward Amazon's offers, leaving non-Amazon listings with fewer sales and little opportunity to build the review base that drives future demand.

These dynamics reveal the strong impact from the interaction between the vertical agreement and Amazon's platform strategy. Brand gating removes competitors, but by itself it does not steer demand: products that lost sellers but were not entered by Amazon saw prices fall and reviews stagnate, indicating that seller exit without algorithmic favoritism generates limited market power. Amazon's Buy Box control, in turn, is far more potent when competition has already been removed by contract. Together, the two mechanisms form a self-reinforcing feedback loop. Amazon uses the Buy Box to channel demand toward its own offers, generating the transactions that build review counts and improve ratings. The brand gating agreement raises barriers to entry, preventing other sellers from competing away this advantage. The resulting ratings and review volume further reinforce Amazon's algorithmic position in future periods. The combination of vertical foreclosure and platform strategy thus does not merely reallocate existing demand: it reshapes the competitive landscape over time, progressively concentrating market share and consumer attention on Amazon's listings in a way that neither the contract nor the algorithm could achieve independently.

7 Welfare Implications

The evidence assembled in Sections 6 and 5 allows us to assess the welfare consequences of the Apple–Amazon agreement across three dimensions: prices, product variety, and quality. On each dimension, the findings point in the same direction, and the favorable quality-based justification offered by Apple and Amazon is not supported by the data.

Price effects. The most direct welfare consequence is the increase in the price paid by consumers. The event-study and DDD estimates document an Apple-specific increase of approximately 16 percent in the Buy Box price in the months immediately following the agreement. This headline figure, however, understates the full extent of the price shift because it is an average over a changing product mix. Section 6 identifies four reinforcing price channels. First, the exit of lower-priced resellers reduced intra-brand competition for each listing, removing the downward competitive pressure that rival sellers had previously exerted on the Buy Box price. Second, the selective removal of older, lower-priced Apple tablets from the platform shifted the composition of available products upward, raising the average price of the consideration set even holding individual listing prices constant. Third, cross-seller price dispersion on the same product narrowed sharply after the agreement, and newly introduced products exhibited the least dispersion of all, indicating a shrinking scope for independent, competitive price-setting even among the sellers who remained. Fourth, Amazon's entry as the dominant Buy Box seller on newly introduced, high-priced models further elevated the effective transaction price, as its algorithmic self-preferencing insulated these listings from competitive bids. These four mechanisms are mutually reinforcing and operate simultaneously.

The distributional implications of this price increase are uneven. The most affordable tier of the Apple tablet market—older, refurbished, or gray-market models sold by third-party resellers at competitive prices—was disproportionately affected. Consumers who relied on these offers to access Apple products at lower price points either paid more, substituted toward lower-quality competing products, or exited the market altogether. The effective price floor for Apple tablets on Amazon rose, compressing the accessible range of the market for price-sensitive buyers.

Variety and choice effects. Beyond prices, the agreement produced a large reduction in the variety of products available to consumers. The number of active Apple tablet listings in treated markets fell by approximately 55 percent relative to baseline, with treated markets going from an average of around 450 distinct ASINs per month to fewer than 200 by early 2019. This reduction is specific to Apple and does not appear in the Samsung data used as the within-country benchmark.

The variety loss has two welfare-relevant dimensions. First, consumers who had a preference for specific legacy models—due to familiarity, compatibility with accessories, or budget constraints—lost access to their preferred options. The agreement did not merely reduce how many sellers offered each product; it eliminated many products from the platform entirely. Second, the composition of the surviving assortment shifted toward newer, more expensive models. As documented in the stayers-versus-entrants analysis, the consideration set consumers faced after the agreement was concentrated at higher price points than before, constraining the set of options available to any consumer whose willingness to pay fell below the new price floor.

Moreover, the sellers removed were not at the competitive fringe: before the policy, exiters collectively held approximately 55–60 percent of iPad buy box time in treated markets, compared to 25–30 percent for stayers and Amazon combined. The reduction in variety was therefore accompanied by the forced exit of the commercially active and buy-box-dominant supply-side participants.

Quality effects and the efficiency defense. Apple and Amazon publicly justified the agreement on quality and brand-integrity grounds: restricting the reseller channel, they argued, would remove unauthorized sellers of potentially counterfeit or substandard products. On this view, the price and variety effects documented above could, in principle, be offset by an improvement in average product quality.

The evidence does not support this defense. The analysis in Section 6 finds no statistically or economically significant difference in product ratings between the products that exited and those that remained. Sentiment analysis of pre-period review text confirms the absence of a quality gap between exiters and stayers: predicted review scores differ by less than 0.07 stars on a five-point scale, and neither a Welch t-test nor a Mann-Whitney test rejects the null of equal quality. Across the full sample, the DDD estimates for average ratings are economically negligible and statistically imprecise, providing no evidence that the surviving or newly introduced products were rated more favorably than those removed.

The seller-level evidence makes the quality defense even harder to sustain. In treated markets, the sellers removed by brand gating were actually better-rated than those who remained. This reversal is specific to treated markets; in control countries, natural seller exit was correlated with lower ratings, as standard quality-screening would predict. Moreover, FBA adoption among removed sellers (50 percent) was similar to that of stayers (53 percent), ruling out the argument that excluded sellers lacked the logistics infrastructure associated with quality fulfillment.

The stayers-versus-entrants comparison reinforces this conclusion from a different angle. Newly introduced products entered at prices approximately 80 percent above those of stayers, yet their average ratings were just slightly above the stayers throughout the post-agreement period, reaching 4.2–4.5 on a five-point scale—the same range as the products they displaced. The price premium commanded by entrants reflects their position in the product lifecycle, not any measurable quality advantage.

Harm to competition among sellers. Beyond consumer welfare, the agreement produced a structural harm to the seller side of the market. The sellers excluded were not low-quality or inexperienced: they had similar months active on iPads, similar iPad share of total storefront (26 percent vs. 23 percent), and better customer ratings. What distinguished them was scale: exiters averaged 5 thousand total storefront ASINs versus 28 thousand for stayers – a nearly six-fold gap. The agreement thus introduced an implicit scale threshold that systematically excluded small, high-quality, iPad-specialized merchants in favor of large incumbents capable of absorbing margin pressure across a broad product portfolio. This concentration of market access among large-scale sellers is a harm to the competitive structure of the marketplace that is analytically distinct from the price and variety effects on consumers, and unlikely to be corrected by market forces once the entry barriers imposed by brand gating are in place.

Information distortions. A further welfare dimension concerns consumer information. Amazon's Buy Box control channeled the majority of purchase transactions to its own listings, and these products then accumulated substantially more reviews than comparable ASINs not entered by Amazon. Since review volume is both a signal of product quality and an input into future purchase decisions, this creates an asymmetry: Amazon-entered products appear more popular and better-reviewed, while alternatives with comparable intrinsic quality but lower transaction volume remain obscure. This informational distortion is self-reinforcing—concentrated demand generates concentrated reviews, which in turn concentrate future demand—and may persist long after the immediate price and variety effects have stabilized.

Summary. The Apple–Amazon agreement produced adverse welfare effects across all three dimensions assessed. Prices rose by approximately 16 percent on the headline measure, with the true impact on the accessible price frontier likely larger given the removal of the most competitive low-priced offers. Product variety fell by around 55 percent, concentrating the consideration set at higher price points. Average quality, as measured by product ratings, showed no improvement that could offset these costs. The mechanism behind these outcomes combines vertical foreclosure through brand gating with Amazon's algorithmic self-preferencing in a mutually reinforcing way: gating removes the competitors who could undercut Amazon, while Buy Box dominance steers demand toward Amazon's higher-priced offers and accelerates the accumulation of reviews that entrench its position further. The evidence is thus inconsistent with a procompetitive or efficiency-enhancing interpretation of the agreement.

8 Conclusion

This paper studies how vertical contracts and platform algorithms interact to reshape competition in digital markets. The 2018 Apple–Amazon brand-gating agreement provides a sharp and well-identified setting: a bilateral contract between two dominant firms that simultaneously restructured seller access to Apple products across multiple major marketplaces and granted Amazon a privileged position as an authorized reseller. Using iPad product-level panel data from Keepa combined with Amazon Reviews text data, and exploiting cross-country variation in treatment together with Samsung as an within-country benchmark, we document the causal effects of this agreement on market outcomes.

Our empirical analysis yields six main findings. First, the agreement triggered a large and persistent contraction in the supply side of the Apple tablet market: the average number of sellers fell by 93 percent,[28] and the number of distinct models available to consumers declined by around 55 percent. This product variety loss was not random—it was concentrated among older,[29] lower-priced models whose independent resellers were removed, while newer, higher-priced models entered or survived. Second, Amazon's Buy Box share for Apple products surged from near zero to above 50 percent in European markets and approached 100 percent in Japan, a gain that is specific to Apple and absent in the Samsung benchmark; this gain fell disproportionately on merchant-fulfilled third-party sellers rather than those fulfilling through Amazon's own network,

[28]the average number of sellers per product fell by approximately 33 percent

[29]This was driven by cohort-based selection that was substantially stronger in treated than control markets, ruling out natural aging alone.

and even Amazon's win rate relative to same-channel FBA sellers exceeds what fulfillment logistics alone would predict. Third, Buy Box prices rose by approximately 16 percent in treated markets relative to controls, but this aggregate figure masks a critical asymmetry: prices rose only for products that Amazon subsequently entered; products not entered by Amazon actually declined in price despite losing the bulk of their seller base. Fourth, the product composition of the consideration set shifted upward: exiters were not lower-quality products—their ratings were statistically indistinguishable from those of stayers—but they were older and more price-competitive. New entrants arrived at prices roughly 80 percent above those of surviving legacy products, at comparable quality, reshaping the effective price floor of the market. Fifth, removed sellers were not marginal operators: exiters were better-rated than stayers in treated markets, had near-identical FBA adoption, and held approximately 55–60 percent of pre-policy Buy Box share. The only difference was scale in scale – exiters averaged 5 thousand storefront ASINs versus 28 for stayers – not quality. Sixth, Amazon's Buy Box dominance generated a self-reinforcing informational asymmetry: channeling transactions to its own listings caused those products to accumulate reviews rapidly, further entrenching Amazon's algorithmic advantage in future periods.

Taken together, these findings point to a *algorithmically mediated demand steering enabled by vertical contracts* mechanism. The brand-gating agreement alone does not generate sustained price increases: removing competing sellers from products that Amazon does not enter leaves prices to fall under residual competition. Brand gating is, however, not inert on its own—by mechanically narrowing the number and diversity of sourcing channels behind each listing, it compresses the dispersion of prices sellers could otherwise offer, removing the competitive slack that Amazon's entry would otherwise have to overcome. The Buy Box algorithm then converts this narrowed channel into higher and more durable prices, rather than being the sole cause: contractual foreclosure supplies the necessary precondition, and algorithmic allocation supplies the sufficient one. Together they convert contractual access into the ability to protect prices and accumulate informational advantages.

These findings speak directly to ongoing policy debates. Apple and Amazon publicly justified the agreement on brand integrity and counterfeit-prevention grounds. Our evidence does not support this defense: exiting products were rated equivalently to surviving ones, sentiment analysis finds no quality differential, and DDD estimates for ratings are economically negligible. The efficiency defense is inconsistent with the data. More broadly, the results suggest that standard antitrust frameworks focused on price coordination or horizontal market division may miss a distinct form of competitive harm in platform markets: the joint foreclosure of resellers and consumers through the complementarity of contractual access and algorithmic demand control. The Digital Markets Act's obligations on self-preferencing and platform neutrality are directly relevant to this mechanism, as is scrutiny of vertical agreements between dominant platforms and premium suppliers. The case illustrates that the line between a legitimate authorized-dealer program and an anticompetitive foreclosure strategy depends critically on whether the platform uses its intermediary position—specifically, its algorithmic demand-allocation tools—to leverage the contractual access it obtains.

Several limitations point to directions for future work. Our analysis is restricted to the short-run window around the agreement's implementation; longer-run effects on innovation, on

Apple's own product strategy, and on inter-brand tablet competition remain to be studied. At the same time, our welfare assessment is qualitative: structural estimation of the demand-side losses from the reduction in variety and the upward shift in prices would provide a cleaner dollar figure of the consumer harm. These extensions offer a natural agenda for future empirical work on the intersection of vertical restraints, algorithmic allocation, and competition policy in digital markets.

References

Aghion, P. and Bolton, P. (1987). Contracts as a barrier to entry. *The American economic review*, pages 388–401.

Autorità Garante della Concorrenza e del Mercato (AGCM) (2023). Apple–amazon proceedings in italy. Competition authority materials related to Apple and Amazon agreements.

Bedre-Defolie, O., Johansen, B. O., and Madio, L. (2024). Design and governance of quality on a digital platform. Technical report, Mimeo.

Casner, B. (2020). Seller curation in platforms. *International Journal of Industrial Organization*, 72:102659.

Chen, N. and Tsai, H.-T. (2024). Steering via algorithmic recommendations. *The RAND Journal of Economics*, 55(4):501–518.

Comisión Nacional de los Mercados y la Competencia (CNMC) (2023). Apple/amazon in spain: Decision and findings. Spanish competition authority decision summarizing the Amazon–Apple agreements and their effects.

Daughety, A. F. and Reinganum, J. F. (1995). Product safety: liability, r&d, and signaling. *The American Economic Review*, pages 1187–1206.

Daughety, A. F. and Reinganum, J. F. (2006). Markets, torts, and social inefficiency. *The RAND Journal of Economics*, 37(2):300–323.

De Chiara, A., Ganuza, J. J., Gómez, F., Manna, E., and Segura-Moreiras, A. (2024). Platform liability with reputational sanctions. *UB Economics–Working Papers, 2024, E24/467*.

De Corniere, A. and Taylor, G. (2014). Integration and Search Engine Bias. *The RAND Journal of Economics*, 45(3):576–597.

De Corniere, A. and Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, 50(4):854–882.

Decarolis, F. and Li, M. (2025). Gatekeepers and Self-Preferencing: Incentives and Welfare Trade-offs in Two-sided Markets. *CEPR Working Paper*.

Decarolis, F., Li, M., and Paternollo, F. (2025). Competition and Defaults in Online Search. *American Economic Journal: Microeconomics (2025)*.

Drugov, M. and Jeon, D.-S. (2019). Vertical integration and algorithm bias. Technical report, mimeo.

Elfenbein, D. W., Fisman, R., and McManus, B. (2015). Market structure, reputation, and the value of quality certification. *American Economic Journal: Microeconomics*, 7(4):83–108.

Farronato, C., Fradkin, A., and MacKay, A. (2023). Self-Preferencing at Amazon: Evidence from Search Rankings. Technical report, AEA Papers and Proceedings.

Gómez-Losada, Á. and Duch-Brown, N. (2019). Competing for amazon's buy box: A machine-learning approach. In *International Conference on Business Information Systems*, pages 445–456. Springer.

Gutiérrez, G. (2022). The Welfare Consequences of Regulating Amazon.

Hagiu, A., Teh, T.-H., and Wright, J. (2022). Should platforms be allowed to sell on their own marketplaces? *The RAND Journal of Economics*, 53(2):297–327.

Hagiu, A. and Wright, J. (2020). Platforms and the exploration of new products. *Management Science*, 66(4):1527–1543.

Hui, X., Saeedi, M., Shen, Z., and Sundaresan, N. (2016). Reputation and regulations: Evidence from ebay. *Management Science*, 62(12):3604–3616.

Inderst, R. and Valletti, T. (2011). Incentives for input foreclosure. *European Economic Review*, 55(6):820–831.

Jiménez Durán, R., Müller, K., and Schwarz, C. (2022). The effect of content moderation on online and offline hate: Evidence from germany's netzdg. *Available at SSRN 4230296*.

José Ganuza, J., Gomez, F., and Robles, M. (2016). Product liability versus reputation. *The Journal of Law, Economics, and Organization*, 32(2):213–241.

Karle, H., Peitz, M., and Reisinger, M. (2020). Segmentation versus agglomeration: Competition between platforms with competitive sellers. *Journal of Political Economy*, 128(6):2329–2374.

Kittaka, Y. and Sato, S. (2022). Dual-Role Platforms and Self-Preferencing: Sequential Search Approach. *Available at SSRN 3736574*.

Landes, W. and Lichtman, D. (2003). Indirect liability for copyright infringement: Napster and beyond. *Journal of economic perspectives*, 17(2):113–124.

Liu, Y., Yildirim, P., and Zhang, Z. J. (2022). Implications of revenue models and technology for content moderation strategies. *Marketing Science*, 41(4):831–847.

Lizzeri, A. (1999). Information revelation and certification intermediaries. *The RAND Journal of Economics*, pages 214–231.

Motta, M. (2023). Self-preferencing and foreclosure in digital markets: Theories of harm for abuse cases. *International Journal of Industrial Organization*, 90:102974.

Padilla, J., Perkins, J., and Piccolo, S. (2022). Self-Preferencing in Markets with Vertically Integrated Gatekeeper Platforms. *The Journal of Industrial Economics*, 70(2):371–395.

Polinsky, A. M. and Shavell, S. (2012). Mandatory versus voluntary disclosure of product risks. *The Journal of Law, Economics, & Organization*, 28(2):360–379.

Raskovich, A. (2007). Retail buyer power through steering. *Economics Letters*, 96(2):221–225.

Raval, D. (2022). Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box. *Working Paper*.

Reimers, I. and Waldfogel, J. (2021). Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review*, 111(6):1944–71.

Rey, P. and Tirole, J. (1986). The logic of vertical restraints. *The American economic review*, pages 921–939.

Rey, P. and Vergé, T. (2004). Bilateral control with vertical contracts. *RAND Journal of Economics*, pages 728–746.

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

U.S. federal court antitrust complaint concerning alleged Apple–Amazon agreement (2023). Amended complaint, floyd v. apple inc. and amazon.com, inc.

Waldfogel, J. (2024). Amazon Self-Preferencing in the Shadow of the Digital Markets Act. Technical report, National Bureau of Economic Research.

Zennyo, Y. (2022). Platform Encroachment and Own-Content Bias. *The Journal of Industrial Economics*, 70(3):684–710.

Zennyo, Y. (2023). Should platforms be held liable for defective third-party goods? *International Economic Review, 2026*.

Zhu, F. and Liu, Q. (2018). Competing with complementors: An empirical look at amazon. com. *Strategic management journal*, 39(10):2618–2642.

A Appendix

A.1 Methodology

Figure A.1: Illustrative event-study path

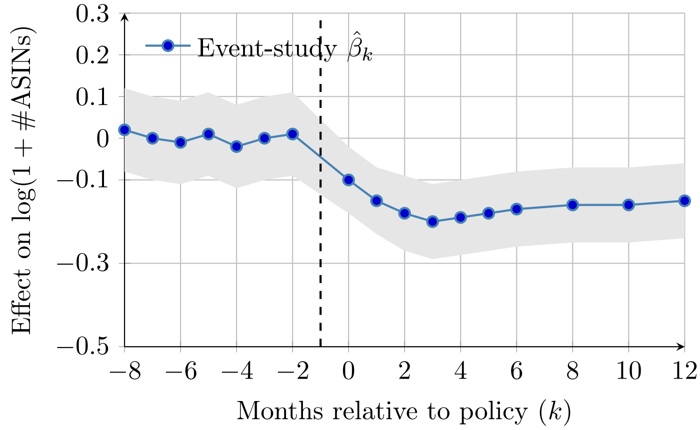

Note: Apple in treated vs. control countries (coefficients $\hat{\beta}_k$ from (1)). Shaded band depicts schematic confidence intervals.

Figure A.2: Illustrative dynamic triple-difference path

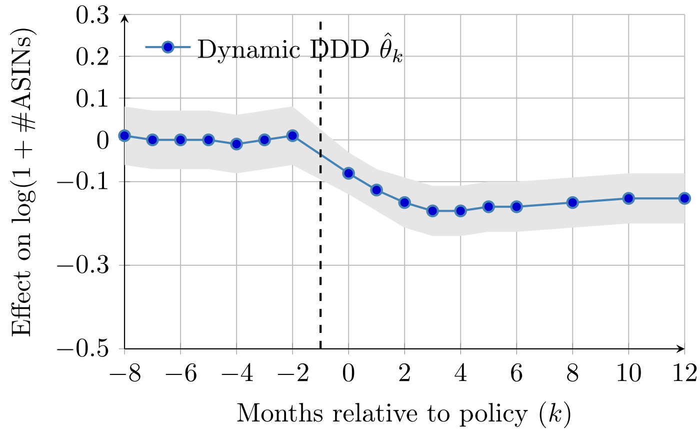

Note: Coefficients $\hat{\theta}_k$ from (2)), contrasting Apple vs. Samsung in treated vs. control countries.

A.2 Results

To begin with, we examine the monthly averages of the number of sellers (including both third-party resellers and Amazon) offering each iPad in treated and control markets, as shown in Figure A.3a. Before the agreement, treated countries averaged about 5 sellers per product, compared with approximately 1.5 in the control group. After the agreement took effect, seller counts in treated markets fell sharply and stabilized at approximately 2, while the control group remained unchanged.

In Figures A.3b and A.3d, we report the corresponding event-study and dynamic DDD estimates, which confirm this pattern. Pre-period coefficients are centered at zero and generally statistically indistinguishable from zero in both specifications, supporting the parallel pre-trends assumption. Post-implementation, the event-study coefficients turn sharply negative, stabilizing near -0.4 log points through months three to nine, implying a roughly **33%** decline in active Apple resellers. The triple-difference estimates, which use Samsung tablets as a within-country benchmark, yield a consistent result. The Apple–Samsung seller gap widens abruptly in treated markets after January 2019 and remains flat in the control group (Figures A.3d and A.3c), which rules out platform-wide trends or common seasonal variation. The regional breakdown in Figure A.3e also shows that the decline is not concentrated in any single country but appears consistently across all treated markets.

Figure A.3: Impact of the Apple–Amazon agreement on the number of Apple tablet sellers per ASIN on Amazon

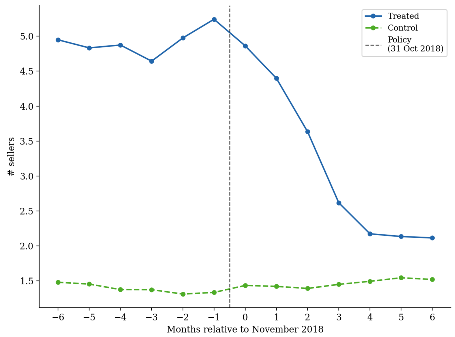

(a) Raw averages: treated vs control

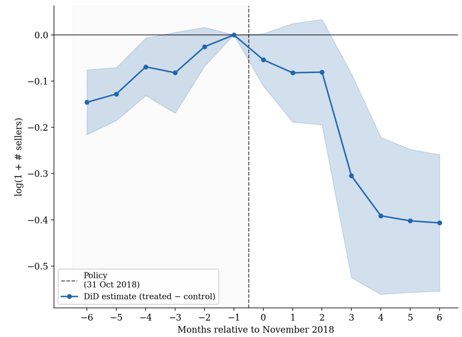

(b) Event Study: iPad sellers

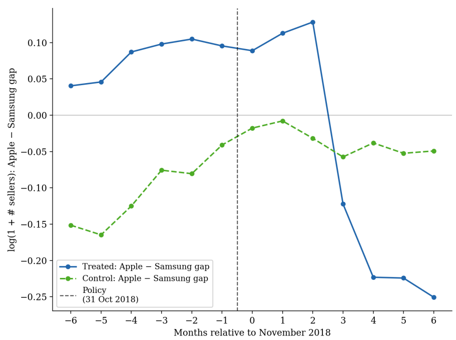

(c) Diagnostic: Apple–Samsung difference

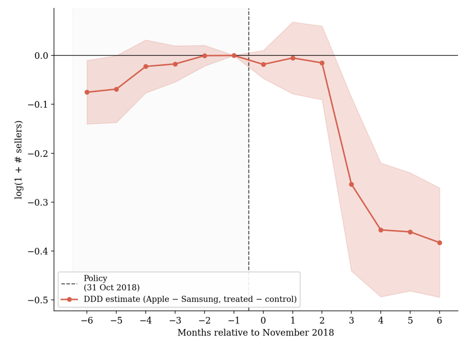

(d) Dynamic Triple-Difference: Apple vs Samsung

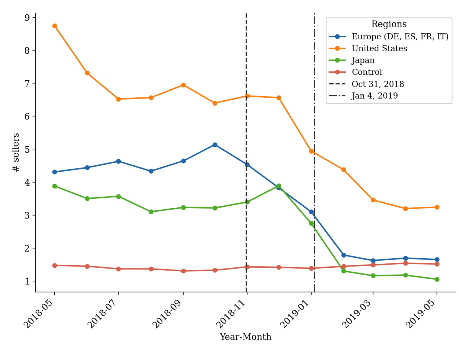

(e) Average monthly number of iPad sellers per ASIN by region

Note: Figure A.3a reports raw averages; Figure A.3b shows the event-study estimates; Figure A.3c shows the raw diagnostic Apple–Samsung difference; Figure A.3d presents the dynamic triple-difference estimates contrasting Apple with Samsung; and Figure A.3e shows the raw averages by region.

Table A.1: Event-Study: Apple Tablets on Amazon (treated vs. control)

	log(1 + avg # sellers)	log(1 + #ASINs)	Amazon BB share	log(1 + BB price)	Rating (1–5)	Seller Rating (0–100)
Pre-policy months (relative to baseline $k = -1$)						
$k = -6$	-0.146^{***}	$+0.237^{**}$	-0.003	-0.105^{***}	$+0.004$	$+4.089^{***}$
	(0.036)	(0.106)	(0.013)	(0.034)	(0.021)	(1.551)
	[Δ%: −13.6]	[Δ%: +26.8]		[Δ%: −10.0]		
$k = -5$	-0.128^{***}	$+0.223^{**}$	-0.004	-0.097^{***}	$+0.019$	$+3.830^{*}$
	(0.029)	(0.092)	(0.008)	(0.033)	(0.028)	(2.007)
	[Δ%: −12.0]	[Δ%: +25.0]		[Δ%: −9.2]		
$k = -4$	-0.069^{**}	$+0.144^{*}$	-0.010^{**}	-0.093^{***}	$+0.040$	$+1.429^{**}$
	(0.032)	(0.078)	(0.005)	(0.035)	(0.029)	(0.696)
	[Δ%: −6.7]	[Δ%: +15.5]		[Δ%: −8.9]		
$k = -3$	-0.082^{*}	$+0.054$	-0.010	-0.034^{***}	-0.010	$+0.731^{***}$
	(0.045)	(0.073)	(0.007)	(0.012)	(0.017)	(0.198)
	[Δ%: −7.9]	[Δ%: +5.6]		[Δ%: −3.3]		
$k = -2$	-0.026	$+0.026$	-0.005	-0.049^{***}	-0.023^{**}	$+0.323$
	(0.021)	(0.046)	(0.004)	(0.017)	(0.011)	(0.228)
	[Δ%: −2.6]	[Δ%: +2.6]		[Δ%: −4.8]		
Post-policy months						
$k = 0$	-0.054^{*}	-0.071^{*}	$+0.014$	$+0.065^{***}$	$+0.080^{***}$	-0.319
	(0.029)	(0.043)	(0.012)	(0.014)	(0.020)	(0.905)
	[Δ%: −5.3]	[Δ%: −6.9]		[Δ%: +6.7]		
$k = 1$	-0.082	-0.109^{**}	$+0.059^{***}$	$+0.131^{***}$	$+0.119^{***}$	-0.315
	(0.054)	(0.045)	(0.021)	(0.029)	(0.021)	(1.536)
	[Δ%: −7.9]	[Δ%: −10.3]		[Δ%: +14.0]		
$k = 2$	-0.080	-0.428^{***}	$+0.218^{***}$	$+0.177^{***}$	$+0.135^{***}$	-1.095
	(0.058)	(0.112)	(0.034)	(0.065)	(0.028)	(1.186)
	[Δ%: −7.7]	[Δ%: −34.8]		[Δ%: +19.4]		
$k = 3$	-0.305^{***}	-0.806^{***}	$+0.335^{***}$	$+0.153^{*}$	$+0.124^{***}$	-1.666
	(0.113)	(0.199)	(0.070)	(0.091)	(0.034)	(1.268)
	[Δ%: −26.3]	[Δ%: −55.3]		[Δ%: +16.5]		
$k = 4$	-0.391^{***}	-0.689^{***}	$+0.415^{***}$	$+0.128^{*}$	$+0.123^{***}$	-2.511
	(0.087)	(0.193)	(0.084)	(0.078)	(0.043)	(1.631)
	[Δ%: −32.4]	[Δ%: −49.8]		[Δ%: +13.7]		
$k = 5$	-0.402^{***}	-0.736^{***}	$+0.427^{***}$	$+0.078$	$+0.138^{***}$	-1.540
	(0.079)	(0.178)	(0.089)	(0.089)	(0.052)	(0.947)
	[Δ%: −33.1]	[Δ%: −52.1]		[Δ%: +8.1]		
$k = 6$	-0.407^{***}	-0.652^{***}	$+0.424^{***}$	$+0.120^{**}$	$+0.136^{***}$	-2.156
	(0.075)	(0.205)	(0.090)	(0.060)	(0.047)	(1.661)
	[Δ%: −33.4]	[Δ%: −47.9]		[Δ%: +12.7]		
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	Country	Country	Country	Country	Country	Country
N	33,136	104	23,194	24,249	27,355	104
R^2	0.096	0.921	0.296	0.131	0.042	0.930

Notes: Coefficients are event-time DiD estimates for treated countries relative to baseline $k = -1$ (month before policy). Standard errors clustered by country in parentheses. Rows labeled Δ% report $100 \times (\exp(\hat{\beta}) - 1)$ for log outcomes (columns 1, 2, and 4). Columns 5 and 6 are estimated in levels. Seller Rating (col. 6) is measured on a 0–100 scale and aggregated to the country×month level. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.2: Triple-Difference: Apple vs. Samsung (treated × post × Apple)

	log(1 + avg # sellers)	log(1 + #ASINs)	Amazon BB share	log(1 + BB price)	Rating (1–5)	Seller Rating (0–100)
DDD estimate	-0.377^{***}	-0.585^{***}	$+0.235^{***}$	-0.154^{*}	-0.076	$+2.330$
	(0.037)	(0.152)	(0.034)	(0.088)	(0.054)	(1.912)
Implied % change	−31.4%	−44.3%	—	−14.3%	—	—
Country×Brand FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	Country	Country	Country	Country	Country	Country
N	692,576	1,715	568,164	552,624	618,962	1,855
R^2	0.014	0.137	0.010	0.002	0.000	0.018

Notes: Each column reports the triple-difference coefficient for treated × post × Apple. Standard errors clustered by country in parentheses. "Implied % change" is $100 \times (\exp(\hat{\beta}) - 1)$ and is shown only for log outcomes (columns 1, 2, and 4). Columns 5 and 6 are in levels. Seller Rating (col. 6) is measured on a 0–100 scale and aggregated to the country×month level. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3: Dynamic Triple-Difference: Apple vs. Samsung (treated × Apple × month)

	log(1 + avg # sellers)	log(1 + #ASINs)	Amazon BB share	log(1 + BB price)	Rating (1–5)	Seller Rating (0–100)
Pre-policy months (relative to baseline $k = -1$)						
$k = -6$	-0.075^{**}	$+0.211^{***}$	-0.082^{***}	-0.038	$+0.025$	$+0.787$
	(0.033)	(0.062)	(0.027)	(0.037)	(0.021)	(0.992)
	[Δ%: −7.2]	[Δ%: +23.5]		[Δ%: −3.7]		
$k = -5$	-0.069^{**}	$+0.195^{***}$	-0.069^{***}	-0.038	$+0.033^{*}$	$+0.870$
	(0.035)	(0.056)	(0.018)	(0.028)	(0.019)	(1.046)
	[Δ%: −6.7]	[Δ%: +21.5]		[Δ%: −3.7]		
$k = -4$	-0.022	$+0.166^{***}$	-0.030^{***}	-0.037	$+0.051^{**}$	$+0.204$
	(0.028)	(0.044)	(0.008)	(0.028)	(0.020)	(0.503)
	[Δ%: −2.2]	[Δ%: +18.0]		[Δ%: −3.6]		
$k = -3$	-0.017	$+0.074^{***}$	-0.029^{***}	-0.010	$+0.030^{*}$	$+0.261$
	(0.019)	(0.022)	(0.009)	(0.013)	(0.016)	(0.206)
	[Δ%: −1.7]	[Δ%: +7.7]		[Δ%: −1.0]		
$k = -2$	-0.000	$+0.036^{***}$	-0.031^{***}	-0.021^{**}	$+0.005$	$+0.244$
	(0.011)	(0.012)	(0.011)	(0.010)	(0.011)	(0.175)
	[Δ%: −0.0]	[Δ%: +3.7]		[Δ%: −2.1]		
Post-policy months						
$k = 0$	-0.018	-0.004	$+0.017$	$+0.052^{***}$	$+0.013$	-0.302
	(0.015)	(0.015)	(0.010)	(0.015)	(0.015)	(0.336)
	[Δ%: −1.8]	[Δ%: −0.4]		[Δ%: +5.3]		
$k = 1$	-0.005	-0.043	$+0.070^{***}$	$+0.102^{***}$	$+0.015$	-0.549
	(0.037)	(0.037)	(0.018)	(0.024)	(0.023)	(0.479)
	[Δ%: −0.5]	[Δ%: −4.2]		[Δ%: +10.7]		
$k = 2$	-0.015	-0.421^{***}	$+0.231^{***}$	$+0.145^{***}$	$+0.030$	-0.795
	(0.038)	(0.095)	(0.035)	(0.056)	(0.029)	(0.530)
	[Δ%: −1.5]	[Δ%: −34.4]		[Δ%: +15.6]		
$k = 3$	-0.263^{***}	-0.836^{***}	$+0.354^{***}$	$+0.155^{**}$	$+0.036$	-1.032^{*}
	(0.091)	(0.159)	(0.071)	(0.071)	(0.031)	(0.576)
	[Δ%: −23.1]	[Δ%: −56.6]		[Δ%: +16.7]		
$k = 4$	-0.357^{***}	-0.715^{***}	$+0.434^{***}$	$+0.127^{**}$	$+0.035$	-1.418^{**}
	(0.070)	(0.141)	(0.082)	(0.064)	(0.036)	(0.675)
	[Δ%: −30.0]	[Δ%: −51.1]		[Δ%: +13.5]		
$k = 5$	-0.361^{***}	-0.723^{***}	$+0.434^{***}$	$+0.107^{*}$	$+0.052$	-1.172^{***}
	(0.062)	(0.144)	(0.091)	(0.055)	(0.044)	(0.434)
	[Δ%: −30.3]	[Δ%: −51.5]		[Δ%: +11.3]		
$k = 6$	-0.383^{***}	-0.730^{***}	$+0.418^{***}$	$+0.107^{**}$	$+0.053$	-1.176^{*}
	(0.057)	(0.148)	(0.089)	(0.054)	(0.041)	(0.698)
	[Δ%: −31.8]	[Δ%: −51.8]		[Δ%: +11.3]		
Country×Brand FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Clusters	Country	Country	Country	Country	Country	Country
N	63,194	208	43,541	44,806	49,212	208
R^2	0.009	0.745	0.086	0.004	0.000	0.111

Notes: Each column reports triple-difference event-time coefficients for treated × Apple at month k, relative to baseline $k = -1$ (month before policy). Standard errors clustered by country in parentheses. Rows labeled Δ% show $100 \times (\exp(\hat{\beta}) - 1)$ and appear only for log outcomes (columns 1, 2, and 4). Columns 5 and 6 are in levels. Seller Rating (col. 6) is measured on a 0–100 scale and aggregated to the country×month level. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.3 Robustness: Minimum New-Condition Price

Figure A.4: Robustness: Impact of the Apple–Amazon agreement on the minimum new-condition price of Apple tablets

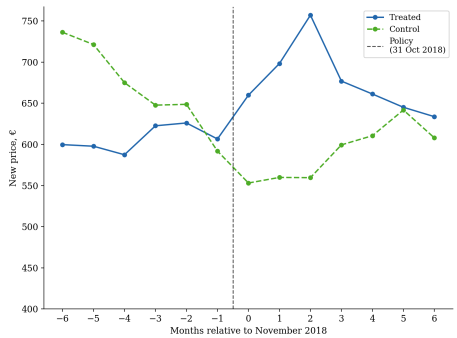

(a) Raw averages: treated vs. control

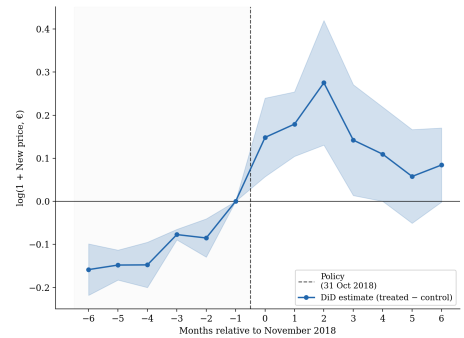

(b) Event Study

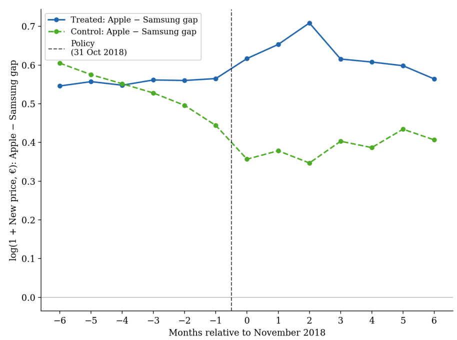

(c) Diagnostic: Apple–Samsung difference

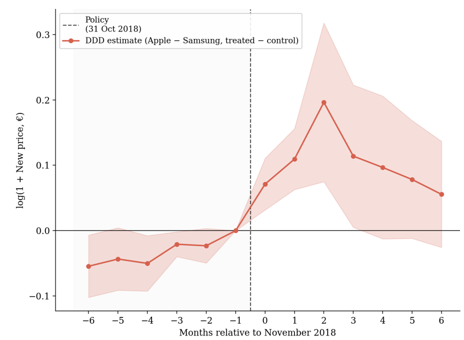

(d) Dynamic Triple-Difference: Apple vs. Samsung

Notes: The outcome is the minimum new-condition listing price (including shipping) in euros, adjusted for inflation. This serves as a robustness check on the Buy Box price results in Figure 6: unlike the Buy Box price, the minimum price is not mechanically affected by the composition of sellers holding the Buy Box. Figure A.4b shows the DiD event-study estimates; Figure A.4a reports raw averages; Figure A.4d presents the dynamic triple-difference estimates contrasting Apple with Samsung; Figure A.4c shows the raw Apple–Samsung diagnostic.