Sherlocking: The Effects of Platform-Owner Entry on the Competitive Behavior of Third-Party Firms*

Benjamin T. Leyden†

December 1, 2025

This draft contains preliminary results; please do not cite.

The latest draft is available here.

I study how third-party firms respond when a platform owner enters its own marketplace by analyzing 23 instances of Apple entering App Store submarkets between 2016 and 2021. Using a staggered difference-in-differences design, I document striking heterogeneity in competitive responses. Effects on entry/exit, monetization, and quality outcomes vary across markets. Many markets show no meaningful responses, particularly for quality metrics like ratings and update frequency. Within affected markets, impacts vary spatially: apps "close" to Apple's offering often experience substantially different effects than more distant apps. Responses also depend on *how* Apple entered: integrated OS features generate different responses on average compared to standalone apps. This heterogeneity—across markets, spatially within markets, and between entry types suggest a more nuanced approach to regulating platform-owner entry may be warranted relative to categorical, platform-wide restrictions.

*I am grateful to Qihong Liu, Chiara Farronato, Devesh Raval, and participants at the 20^{th} Annual International Industrial Organization Conference, the NYU Economics of Strategy Workshop, the NUS Economics of Platforms Workshop, ASSA 2023, the US DOJ Economic Analysis Group, the Cornell Strategy and Economics of Digital Markets Conference, HKUST workshop, and the 2025 CESifo Economics of Digitization Conference for helpful comments. I thank Yujie Feng for excellent research assistance, and I acknowledge the Cornell Center for Social Sciences for financial support.

†Cornell University; CESifo; leyden@cornell.edu

1 Introduction

Platform-owner entry into digital marketplaces has become a central issue in antitrust policy and the regulation of digital markets(De Chant, 2022; Warren, 2019). The practice—where marketplace operators compete with third-party sellers in their own platforms—has prompted significant regulatory responses worldwide. The European Union's Digital Markets Act, enacted in 2022, explicitly restricts "gatekeepers" from engaging in "any form of differentiated or preferential treatment...in favour of products or services it offers itself" (EUR-Lex, 2022). In the United States, the previously proposed American Innovation and Choice Online Act would have similarly banned platforms from using "nonpublic data...to offer...products or services...that compete...with products or services offered by business users" (Congress.gov, 2022). These blanket restrictions reflect growing concerns that platform owners exploit their unique position to discriminate against competitors, leveraging privileged information and control over marketplace infrastructure to advantage their own offerings (Crémer, de Montjoye, and Schweitzer, 2019; Scott Morton et al., 2019).

Yet the economic effects of platform-owner entry remain theoretically ambiguous and empirically contested (Anderson and Bedre Defolie, 2024; Hagiu, Teh, and Wright, 2022). This "Sherlocking" phenomenon—named after Apple's incorporation of the third-party search app *Watson*'s functionality into its own *Sherlock* product—exemplifies developers' concerns about first-party entry, where platform features can suddenly render entire categories of third-party apps obsolete. But, while entry introduces a typically well-resourced competitor that may engage in self-preferencing or restrict access to essential features, it can also generate positive spillovers through technology diffusion, market expansion, and complementary infrastructure investments. The debate over appropriate regulatory responses—ranging from outright bans to conduct remedies to laissez-faire approaches—hinges critically on understanding the magnitude, heterogeneity, and mechanisms underlying these effects.

In the context of mobile software platforms, platform-owner entry takes two distinct forms that have received insufficient attention in both academic literature and policy discussions. Apple's entry into submarkets within its App Store marketplace occurs through either *standalone entry*—developing and releasing a new separately installable application—or *integrated entry*—incorporating competing functionality directly into the operating system.

Standalone entry involves Apple creating a new app that users can discover and choose to download from the App Store, competing alongside third-party offerings. Examples include Measure (an augmented reality measurement tool), Translate (language translation), and Magnifier (visual assistance). These apps appear in search results, have product pages with descriptions and ratings, and can be uninstalled by users—though some may come pre-installed by default. The competitive dynamic is relatively transparent: consumers actively choose whether to adopt Apple's offering over alternatives.

Integrated entry, by contrast, involves Apple integrating functionality directly into iOS that automatically becomes available to all users upon system updates. Examples include health tracking features integrated into the Health app, weather widget improvements, enhanced Maps functionality, and camera feature additions. This functionality cannot be separately uninstalled, often comes bundled with new APIs that third-party developers can leverage, and represents a more fundamental alteration of the platform's capabilities. Users receive these features automatically, without making an active adoption or use decision.

This distinction matters for competitive dynamics. Integrated entry creates unavoidable competition—third-party apps must compete with functionality that every user already possesses. It may also generate stronger technology spillovers, as OS-level integration often introduces new hardware capabilities or software frameworks that benefit the broader developer ecosystem.

This paper provides comprehensive evidence on platform-owner entry effects by studying 23 instances of Apple entering submarkets in its App Store between March 2016 and September 2021. The analysis encompasses 14 standalone app releases and 9 integrated OS integrations, enabling systematic comparison across entry types and market characteristics.[1]

I make three key contributions to the literature on platform competition and regulation. First, the scale of analysis—examining 24 entry events compared to the handful studied in prior work—enables systematic investigation of heterogeneous treatment effects across markets.[2] This reveals substantial variation in competitive responses that single-market studies cannot detect. Second, I develop a novel empirical approach using continuous distance measures derived from S-BERT embeddings to capture within-market heterogeneity based on competitive proximity. This method-

[1] Work is ongoing to integrate additional markets into this analysis.
[2] Work is ongoing to incorporate additional examples of Apple entry.

ology recognizes that not all apps in an affected market are equally "treated" by platform entry. Third, the analysis documents multi-dimensional heterogeneity in first-party entry effects—both across markets and within markets by competitive distance—that fundamentally challenges one-size-fits-all regulatory approaches.

The empirical strategy leverages a staggered difference-in-differences design, using comparable Android apps and not-yet-treated iOS apps as a control group for iOS apps affected by Apple's entry. Market definitions rely on semantic similarity of app descriptions, with apps exceeding a similarity threshold considered part of the competitive market around each entrant.

The analysis reveals striking heterogeneity that challenges both pro-platform entry and anti-platform entry narratives. Perhaps most surprisingly, a substantial proportion of markets experience no statistically significant effects from Apple's entry in many of the outcomes considered in this analysis. This prevalence of null effects, particularly pronounced for quality metrics, contradicts claims of either universal harm from platform competition.

When effects do occur, they vary across markets and space. Apps close to Apple's offering in semantic space experience substantial, often negative impacts, while more distant competitors remain largely unaffected. This creates a limited "competitive radius" around each entrant which can vary in size by market and outcome.

Entry type proves crucial: integrated OS features generate systematically larger effects than standalone app releases. Price increases are more than twice as large for integrated entry ($0.18) versus standalone ($0.08), while the proportion of free apps declines by 5.2 percentage points for integrated versus only 2.4 percentage points for standalone.

These patterns of heterogeneity—across markets, within markets by distance, and between entry types—suggest that platform-owner entry effects depend critically on contextual factors that current regulatory proposals largely ignore. The evidence strongly supports targeted, context-specific oversight rather than the blanket restrictions embodied in recent legislation.

This paper contributes to several strands of literature on platform competition and digital markets. Most directly, it extends empirical work on platform-owner entry effects. Wen and Zhu (2019) study Google's entry into Android submarkets, finding evidence of reduced innovation and increased prices among affected developers. Zhu and Liu (2018) document that Amazon targets successful product categories and that its entry increases demand but discourages third-party investment.

Foerderer et al. (2018) find that Google's entry into photo management increased both demand and third-party innovation. My analysis expands this literature by examining a much larger set of entry events, enabling systematic analysis of heterogeneous effects.

A related literature examines specific mechanisms of platform advantage. Studies of self-preferencing document that platforms favor their own products in search and recommendations, with mixed welfare implications (Teng, 2023; Lee and Musolff, 2023; Lam, 2023; Raval, 2023). Work on "insider imitation" explores how platforms leverage proprietary data to select entry markets (Madsen and Vellodi, 2025; Hagiu, Teh, and Wright, 2022). While I do not isolate specific mechanisms, my reduced-form estimates capture the net effect of all channels through which platform entry affects competition.

The paper also relates to the "platforms as regulators" literature, which recognizes that platform owners shape competitive dynamics through rules, design choices, and technical standards (Boudreau and Hagiu, 2011). Recent work shows how platform policies on privacy (Kesler, 2022; Li and Tsai, 2022; Johnson et al., 2025), ratings systems (Jabr et al., 2020; Leyden, 2025), and product categorization (Ershov, 2024) affect within-platform competition. My findings suggest that entry decisions represent another regulatory lever, with effects varying by how entry is implemented.

Methodologically, the paper advances techniques for defining and measuring competition in digital markets. Building on Leyden (2023), I use transformer-based language models to construct product spaces from text descriptions, enabling more nuanced market definitions than traditional classification systems. The continuous distance measure allows for gradient effects within markets, recognizing that competitive pressure varies with product differentiation.

The remainder of the paper proceeds as follows. Section 2 describes the data sources, market definition methodology, and sample construction. Section 3 presents the empirical analysis, beginning with identification strategy, then examining market dynamics (Section 3.1: entry, exit, and cohort composition) and incumbent app responses (Section 3.2: monetization and quality outcomes). Section 4 explores mechanisms, focusing on how entry type (standalone versus integrated) shapes competitive responses. Section 5 discusses policy implications and directions for future research.

2 Data

2.1 Data Sources

I use a monthly panel dataset of all apps available in the US App Store and Google Play Store provided by the data analytics company AppMonsta. The full sample runs from January 2015 to June 2022 and consists of 14.4 million unique apps across both platforms. For each app-month, I observe product characteristics including name, price, average rating, rating count, and whether the app was updated that month. The data also includes developer-written text descriptions that serve as apps' primary marketing content and provide a reliable indication of product functionality and purpose. Additionally, I use daily ranking lists for product sales and revenue to construct sample restrictions focused on actively competing apps. Key variables for the analysis include monetization outcomes: Price, Free (indicator for price = \$0), In-App Purchases (IAP) availability, and Price-Paid (price conditional on price > \$0); quality outcomes: Average Rating (1-5 stars), Update indicator, and Rating count (a proxy for consumer demand); and entry/exit outcomes: Market-level counts of new app entries and exits, transformed using inverse hyperbolic sine for regression analysis.

2.2 Market Definition

Given the millions of products observed over the sample period, identifying which third-party apps are affected by platform-owner entry requires precise market definitions. Platform-provided categories (e.g., Games, Productivity) are too broad—the Productivity category alone encompasses password managers, shift scheduling apps, document editors, and many other distinct product types.

I define markets using semantic similarity of product descriptions, leveraging recent advances in natural language processing. Specifically, I apply Sentence-BERT (SBERT), a transformer-based neural network model, to map text descriptions into a 384-dimensional vector space where semantic similarity corresponds to geometric proximity (Reimers and Gurevych, 2019). The model places product vectors such that functionally similar apps—regardless of specific terminology—cluster together.

For any pair of apps i and j with description vectors $desc_i$, $desc_j$, I calculate the cosine similarity

$$d^{\cos}(desc_i, desc_j) = \frac{desc_i^\top desc_j}{\|desc_i\|_2 \, \|desc_j\|_2} \in [-1, 1], \tag{1}$$

where $\|\cdot\|_2$ denotes the Euclidean norm. $d^{\cos}(desc_i, desc_j) = 1$ indicates identical descriptions.

In some of the regression analysis that follows rescale this measure to the unit interval using the "scaled similarity" measure:

$$\tilde{d}^{\cos}(desc_i, desc_j) = \frac{d^{\cos}(desc_i, desc_j) + 1}{2} \in [0, 1]. \tag{2}$$

Given a first-party entrant a with description vector $desc_a$, I define the market J_a^θ as all apps j satisfying:

$$J_a^\theta = \{j \in J^{\text{iOS}} \cup J^{\text{Android}} : d^{\cos}(\text{desc}_a, \text{desc}_j) \geq \theta\} \tag{3}$$

The parameter θ governs market size: higher values impose stricter similarity requirements, yielding smaller, more tightly defined markets, while lower values produce larger markets that include more distant competitors. I use $\theta = 0.6$, which defines tighter, more focused competitive sets than alternative thresholds. This ensures the treated sample consists of apps that are meaningfully similar to the entrant, reducing concerns about including tangentially related products that might bias estimates toward zero. The threshold balances two considerations: setting θ too high yields markets that are too narrow and may miss relevant competitors, while setting it too low includes apps that are not meaningful substitutes.

Each app i in market a thus has a continuous distance measure $d_i \in [\theta, 1]$ from the entrant, enabling analysis of how treatment effects vary with competitive proximity. This continuous treatment approach, developed in Section 3, represents a methodological advance over binary in/out market definitions.

2.3 Entry Events

I study 23 instances of Apple entering submarkets within its App Store ecosystem between March 2016 and September 2021.[3] The manner in which Apple implements entry varies substantially

[3] Work is ongoing to add additional instances of Apple competing in its marketplace.

across these events and has important implications for competitive dynamics. These events divide into two categories based on how functionality is delivered to users. In the first category, *standalone entry*, Apple releases a new separately installable application that users can discover and download from the App Store; my sample includes 14 such events. These apps compete directly alongside third-party offerings with their own product pages, ratings, and the ability to be uninstalled, though some may come pre-installed on devices by default. Examples include Schoolwork (classroom management), Breathe (meditation), and Translate (language translation). In the second category, *integrated entry*, Apple integrates competing functionality directly into iOS through operating system updates; my sample includes 9 such events. This functionality becomes automatically available to all users, cannot be separately uninstalled, and often introduces new APIs that third-party developers can leverage. Examples include Health app enhancements, Apple Music features, camera improvements, and Maps functionality updates.

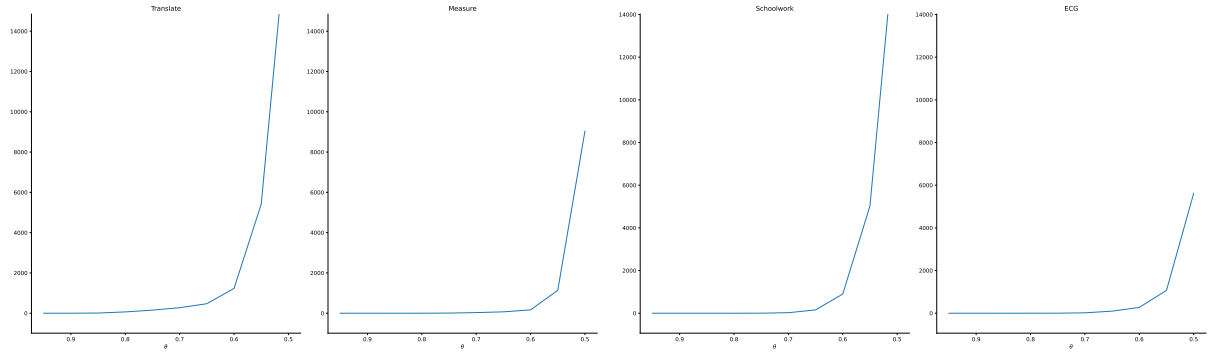
Table 1: Complete List of Apple Entry Events

Market	Entry Date	Type	iOS Apps	Android Apps
Schoolwork	2016-03	Standalone	227	39
Breathe	2016-09	Standalone	137	48
Clock 10.0	2016-09	Integrated	72	80
Health 10.0	2016-09	Integrated	133	39
Homekit 10.0	2016-09	Integrated	113	84
Camera 11.0	2017-09	Integrated	242	187
Development 11.0	2017-09	Integrated	139	15
Notes 11.0	2017-09	Integrated	560	151
Animoji	2017-11	Standalone	203	81
Health 11.3	2018-04	Integrated	98	31
Ibooks 11.4	2018-06	Integrated	118	56
Measure	2018-09	Standalone	71	19
Walkie Talkie	2018-09	Standalone	89	34
Apple Music 13.0.2	2019-09	Integrated	243	56
Ecg	2019-09	Standalone	46	20
Radio	2019-09	Standalone	16	8
Reality Composer	2019-09	Standalone	133	20
Apple Research	2019-11	Standalone	27	7
Sleep	2020-09	Standalone	173	71
Translate	2020-09	Standalone	432	184
Blood Oxygen	2020-10	Standalone	20	7
Airtags	2021-04	Standalone	89	35
Magnifier	2021-09	Standalone	74	22

Notes: App counts based on $\theta = 0.6$ threshold, apps ranked 200+ for 10%+ days.

Table 1 provides the complete list of all 23 entry events with their dates and types. This

Figure 1: Market Size Heterogeneity: Cumulative App Counts by Similarity Threshold



Note: Each panel plots the cumulative number of apps with cosine similarity to the first-party entrant at or above threshold θ. The horizontal axis runs from $\theta = 1.0$ (far left, the entrant's location) to $\theta = 0.5$ (far right). As θ decreases moving rightward, more apps are included in the market.

comprehensive set of entry events, spanning diverse product categories and entry modes, enables systematic analysis of heterogeneous platform-owner entry effects.

2.4 Sample Construction

Market definitions yield initial pools of potentially affected apps for each entry event. I limit analysis to apps that ranked 200th or better for at least 10% of days they were available. This excludes hobby projects, abandoned apps, and experimental products that are not actively competing during the sample period. For each market-entry event, I construct an 18 month balanced panel centered around the first-party entry date.

The resulting dataset contains 76,230 app-months of observations across 23 markets, with market sizes varying substantially based on product space density. Figure 1 illustrates this heterogeneity across four representative markets. Each panel shows how many apps would be included in a market as the similarity threshold θ varies from 0.5 to 1.0. At the baseline threshold of $\theta = 0.6$, prior to imposing the ranking threshold described above, market sizes range from approximately 200 apps in specialized markets (Measure, ECG) to over 1,000 apps in broader product categories (Translate, Schoolwork). The sharp increases visible at $\theta < 0.6$ demonstrate that loosening the threshold would dramatically expand markets to include many tangentially-related apps, motivating our choice of a conservative threshold that focuses on close competitors.

I use similar apps on the Google Play Store as a control group for App Store apps affected by

Apple's entry.[4] Android represents the other major mobile platform with a comparable developer ecosystem and user base. Markets are defined using the same SBERT process across platforms, ensuring that treated iOS apps and control Android apps operate in the same product spaces. The use of Android as a control is particularly compelling because Google does not systematically enter the same markets at the same times as Apple, and both platforms experienced similar technological and demand shocks during the sample period.

2.5 Summary Statistics

Tables 2 and 3 present summary statistics averaged across the 18-month panels (9 months pre-entry, 9 months post-entry) for each of the 23 markets. These tables reveal substantial heterogeneity across markets along multiple dimensions.

Table 2 shows market-level entry and exit patterns. Entry and exit rates vary dramatically across markets: stable markets like Apple Research and Measure average fewer than 0.5 new apps per month, while dynamic markets like News 9.0 and Apple Music 8.4 experience 4–8 entries and exits monthly. Entry and exit are highly correlated within markets—markets with high entry also exhibit high exit, suggesting competitive churn rather than net growth. Across all markets, entry and exit rates are nearly balanced (1.80 entries versus 1.79 exits per month on average), indicating that most markets maintain relatively stable populations despite substantial turnover.

Table 3 presents app-level summary statistics for the 10,119 unique apps comprising 182,142 app-month observations. Monetization strategies vary dramatically across markets. Six markets consist entirely of free apps (Accessibility 10.0, Breathe, Clock 10.0, Health 10.0, Homekit 10.0, and Schoolwork), while others are predominantly paid: Measure (77% paid), ECG (73% paid), and Camera Remote (71% paid). Overall, 70% of apps are free. Among paid apps, average prices vary from \$1.57 (Animoji) to \$6.29 (Translate 14.0), with an overall average of \$4.24. In-app purchase adoption ranges from 23–25% in health and education markets to 79% in the Translate market, averaging 51% across all markets.

Market sizes vary substantially: small markets contain as few as 20–30 apps (Blood Oxygen, Radio, Apple Research), while large markets contain hundreds of apps (Camera 11.0, Notes 11.0,

[4] Android IAP data is only available starting September 2016. As a result, 5 markets (Schoolwork, Breathe, Homekit 10.0, Clock 10.0, and Health 10.0) are excluded from all IAP analyses because they entered before sufficient pre-period data was available.

	IHS Entry	IHS Exit	Raw Entry Count	Raw Exit Count
Airtags	0.67	0.63	0.97	0.91
Animoji	1.42	1.60	3.39	3.56
Apple Music 13.0 (2)	0.93	1.04	1.57	1.81
Apple Music 8.4	2.45	2.69	10.47	10.32
Apple Research	0.16	0.16	0.20	0.19
Blood Oxygen	0.31	0.28	0.41	0.36
Breathe	0.79	0.57	1.24	0.85
Camera 11.0	1.55	1.50	3.22	2.74
Clock 10.0	0.64	0.55	0.91	0.75
Development 11.0	0.43	0.43	0.57	0.62
Ecg	0.27	0.23	0.32	0.30
Health 10.0	0.58	0.43	0.82	0.56
Health 11.3	0.55	0.42	0.78	0.54
Health 8.2	2.09	2.10	6.48	5.07
Health 9.0	0.59	0.52	0.81	0.71
Homekit 10.0	1.07	0.72	1.77	1.05
Ibooks 11.4	0.41	0.42	0.54	0.58
Ibooks 8.4	1.33	1.37	2.60	2.73
Magnifier	0.42	0.41	0.56	0.57
Mail 9.0	0.58	0.73	0.85	1.11
Maps 9.0	2.52	3.02	10.95	12.21
Measure	0.32	0.27	0.41	0.34
News 9.0	1.89	2.25	5.12	6.60
Notes 11.0	1.74	1.75	4.68	4.02
Radio	0.26	0.28	0.35	0.36
Reality Composer	0.65	0.54	1.03	0.71
Schoolwork	0.65	0.67	0.97	0.97
Sleep	0.77	0.66	1.19	0.94
Spotlight Search 9.0	0.90	1.09	1.67	1.97
Translate	1.40	1.54	2.73	3.10
Walkie Talkie	0.85	0.91	1.42	1.48
All Markets	0.94	0.96	2.23	2.19

Table 2: Market-Level Entry/Exit Summary Statistics (Restricted Sample, EntryExit Dataset)

	Price	Price-Paid	Free	IAP	Update	Avg. Rating	IHS # Ratings	# Apps	N
Airtags	1.31	3.98	0.67	0.58	0.25	3.85	7.25	124	2232
Animoji	0.49	1.59	0.69	0.33	0.13	3.84	6.19	284	5112
Apple Music 13.0 (2)	2.35	5.00	0.53	0.32	0.27	3.96	7.23	299	5382
Apple Research	1.80	5.31	0.66	0.36	0.28	4.07	5.64	34	612
Blood Oxygen	1.09	2.27	0.52	0.28	0.16	3.75	5.99	27	486
Breathe	0.00		1.00	0.30	0.14	4.19	5.55	185	3330
Camera 11.0	0.98	4.31	0.77	0.34	0.26	3.97	6.52	429	7722
Clock 10.0	0.00		1.00	0.43	0.19	4.29	7.62	152	2736
Development 11.0	1.90	5.96	0.68	0.34	0.25	3.81	6.07	154	2772
Ecg	3.10	4.35	0.29	0.29	0.14	3.80	5.44	66	1188
Health 10.0	0.00		1.00	0.19	0.28	3.69	5.65	172	3096
Health 11.3	0.85	3.40	0.75	0.19	0.31	3.70	6.03	129	2322
Homekit 10.0	0.00		1.00	0.29	0.38	3.61	6.60	197	3546
Ibooks 11.4	1.81	4.40	0.59	0.23	0.24	3.89	6.58	174	3132
Magnifier	1.47	3.06	0.52	0.38	0.16	3.99	7.95	96	1728
Measure	3.01	3.92	0.23	0.26	0.18	3.97	4.94	90	1620
Notes 11.0	1.22	4.52	0.73	0.41	0.28	4.08	7.22	711	12798
Radio	1.23	2.69	0.54	0.33	0.28	3.90	7.47	24	432
Reality Composer	2.70	4.25	0.36	0.32	0.18	3.96	6.34	153	2754
Schoolwork	0.00		1.00	0.30	0.27	3.94	6.29	266	4788
Sleep	1.40	2.98	0.53	0.44	0.26	4.22	7.77	244	4392
Translate	1.91	4.72	0.59	0.58	0.24	4.21	7.69	616	11088
Walkie Talkie	1.32	3.19	0.59	0.30	0.29	3.56	6.62	123	2214
All Markets	1.23	4.11	0.70	0.38	0.24	3.98	6.79	4749	85482

Table 3: App-Level Summary Statistics (Restricted Sample, App-Base Dataset)

Translate). Update frequencies average 24% (averaging 2-3 updates per year), ranging from 11% in Animoji to 38% in Homekit 10.0.

This heterogeneity in market structure, competitive intensity, and business models motivates the heterogeneous treatment effect analysis in Section 3.2. Markets differ not only in baseline characteristics but also, as I show below, in their responses to platform-owner entry.

3 Empirical Analysis

I estimate the effects of platform-owner entry on third-party developers using a staggered difference-in-differences design that exploits variation in the timing of Apple's entry across 23 submarkets, comparing iOS apps to similar apps on Android's Google Play Store as controls. The analysis examines two dimensions of competitive response: market dynamics (entry, exit, and the composition of entering cohorts) and incumbent responses (monetization and quality outcomes).

The fundamental identification challenge is that we cannot observe counterfactual outcomes—what would have happened to affected apps absent Apple's entry. The identifying assumption is

that absent Apple's entry, iOS and Android apps in the same product space would have followed similar outcome trajectories. The credibility of this design relies on two key institutional features: (1) Google does not systematically enter the same markets at the same times as Apple, avoiding contamination of the control group, and (2) both platforms experience similar technological and demand shocks during the sample period. Requirement (1) holds for all first-party entrants considered here, and (2) is supported by the fact that both Apple and Google's platforms are frontier mobile operating systems, which closely mirror each other in the technological hardware these platforms operate on, and which both compete for largely the same consumers in the mobile phone market. To support this assumption, I present evidence regarding pre-trends through event study specifications throughout my analysis.

A potential threat to identification is spillover effects among multi-homing apps—products available on both platforms. If an iOS app adjusts its strategy in response to Apple's entry, the Android version might also change, contaminating the control group. However, platform-specific development teams, different technology stacks, and distinct user bases often lead to independent strategies across platforms. This specification includes all apps regardless of multi-homing status, maximizing sample size while acknowledging this potential limitation.[5]

3.1 Market Dynamics

I first examine how Apple's entry affects market structure through changes in the number of third-party apps entering and exiting markets, as well as the characteristics of new entrants. To understand the effect of first-party entry on third-party entry and exit, I use a staggered difference-in-differences design, estimated using the Sun and Abraham (2021) interaction-weighted estimator. In Appendix A and the event study plots below, I also report TWFE and ETWFE results for comparison. The outcome variables are market-level counts of new app entries and app exits, transformed using the inverse hyperbolic sine (IHS) function to handle zeros while maintaining interpretability similar to logarithms. Standard errors are clustered at the market-platform level.

Table 4 presents difference-in-differences estimates for entry and exit. Notably, I find no statistically significant effect on third-party entry: an imprecisely estimated 4.5% increase.

[5] Forthcoming analysis will reproduce the results of this paper using only single-homing apps. Preliminary evidence from this exercise is consistent with the results presented here.

The exit regression shows evidence of decline following Apple's entry with the estimate indicating an 18.2% decline in exit, which is also not statistically significant. This evidence that exit is unaffected, or potentially falls, suggests that incumbent apps may be "locked in" once Apple enters, potentially due to sunk development costs or strategic value of maintaining presence despite increased competition. It may also reflect a long-run strategy to retain users as those users become accustomed to the product category and seek alternative or more fully featured applications. Additionally, it is inexpensive for most developers to keep their products in the store even while abandoning active development.I consider this further in Section 3.2.2 by estimating the effect of first-party entry on apps' likelihood of updating.

Table 4: Entry and Exit

	Entry Count (IHS)	Exit Count (IHS)
ATT	0.0452	-0.1820
	(0.1091)	(0.1232)
Baseline Mean	1.192	0.8998
N	1,064	1,064

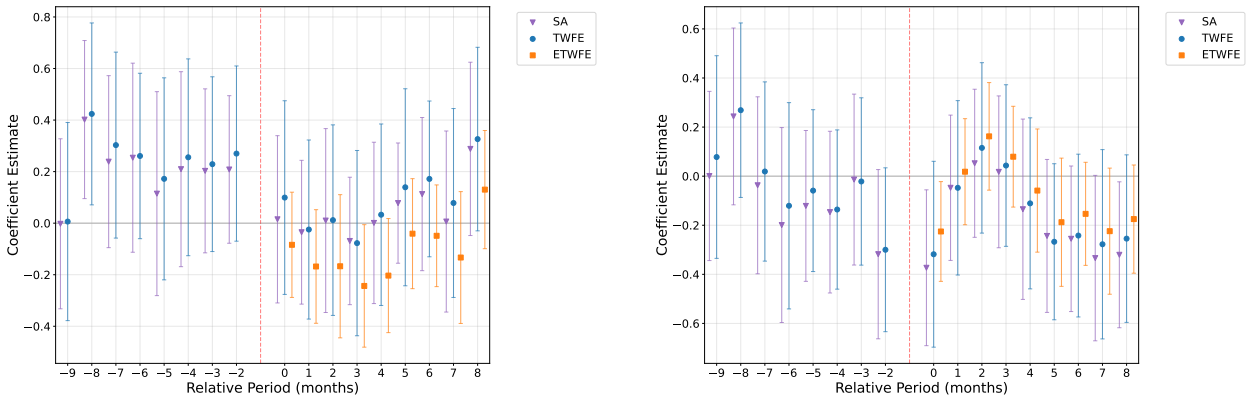
Notes: Sun & Abraham (2021) interaction-weighted DiD estimates at the platform-market-month level. Standard errors clustered at the platform-market level are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Figure 2 presents event study estimates for entry and exit. Pre-entry coefficients for both outcomes hover near zero and show no systematic trends, supporting the parallel trends assumption. Entry effects emerge gradually post-treatment, with coefficients becoming positive four months after entry but remaining imprecise. Exit shows an immediate drop, recovery, and then an additional drop that suggests some developers may choose to exit immediately, while others wait, perhaps to learn more about the impact of Apple's entry.

3.1.1 Composition of Entering Cohorts

In addition to affecting entry and exit counts, platform-owner entry may alter the composition of new entrants. I examine how entering cohorts differ before versus after Apple's entry using a cross-sectional difference-in-differences design. Each entering app is observed once at entry, creating a

Figure 2: Entry and Exit Dynamics



(a) Entry Count (IHS)

(b) Exit Count (IHS)

Note: Event study estimates showing dynamic treatment effects relative to one month before entry. SA (Sun and Abraham (2021)) is the primary estimator; TWFE and ETWFE (Wooldridge (2021)) shown for robustness. Analysis is at the platform-market-month level. Standard errors are clustered at the platform-market level. Error bars represent 95% confidence intervals.

cross-sectional dataset rather than a panel. I estimate:

$$y_i = \alpha \cdot \text{Post}_{at} \times \text{iOS}_i + \delta_{pa} + \eta_t + \varepsilon_i \tag{4}$$

where y_i represents entry characteristics for app i, Post_{at} indicates the app entered after Apple's entry in market a, δ_{pa} are platform-market fixed effects, and η_t are month fixed effects. The coefficient α captures differences in entry characteristics between iOS and Android apps, comparing those entering before versus after Apple's entry. Standard errors are clustered at the market-platform level.

Because each app is observed only once—at its market entry date—this analysis constitutes a repeated cross-section rather than a panel. Estimators like Sun & Abraham that rely on within-unit variation over time are not applicable to this data structure. Accordingly, I estimate Equation (4) using OLS.

Table 5 presents the results from estimating the cohort entry characteristics model Equation (4).[6] The results reveal systematic changes in the composition of entering apps following Apple's entry. While overall prices do not fall, among paid apps, prices decrease substantially—

[6]Event study estimates showing how the iOS-Android gap in entry characteristics evolves relative to Apple's entry timing appear in Appendix B.

new iOS entrants charge approximately $1.21 less than comparable Android entrants, reversing the pattern observed for incumbents. This suggests that developers entering post-treatment may compete more aggressively on price, perhaps recognizing the difficulty of competing with Apple's offerings, which are always free to use, at premium price points. The probability of offering a free app and in-app purchase adoption show no evidence of changes, indicating that monetization strategy choices at entry are not substantially altered by Apple's presence.

Quality metrics at entry show notable changes. Ratings for new iOS entrants are approximately 0.44 stars *lower* than for Android entrants following Apple's entry, suggesting either negative selection (lower-quality developers entering) or that the presence of Apple's product raises the bar against which users evaluate third-party alternatives. Interestingly, despite lower ratings, new iOS entrants receive *more* user engagement as measured by rating counts, possibly reflecting increased category visibility that Apple's entry brings to the market.

Table 5: Cohort ATT Estimates

	Price	Price (Paid Apps)	Free App	In-App Purchases	Rating	Rating Count (IHS)	Similarity
ATT	0.0306	-1.208***	0.0188	0.1422	-0.4421***	0.8417**	0.0124*
	(0.0463)	(0.3891)	(0.0307)	(0.1156)	(0.1580)	(0.3669)	(0.0067)
Baseline Mean	0.1569	3.042	0.9484	0.2932	4.195	2.232	0.6475
N	2,752	181	2,752	932	1,530	2,752	2,752

Notes: Cross-sectional DiD estimates for entering cohorts. Each app observed once at entry. Standard errors clustered at market-platform level shown in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Finally, I estimate the effect of Apple's entry on the initial position of new entrants (Similarity). Following Apple's entry into a market, subsequent entry shifts *closer* to Apple's product. Rather than differentiating away from Apple's offering, new developers appear drawn toward the product space Apple has occupied. This clustering suggests that Apple's entry may serve a coordinating function within its ecosystem. Several mechanisms could drive this pattern: Apple's entry may validate consumer demand for a feature category, reducing uncertainty for prospective developers; new APIs or frameworks released alongside Apple's product may lower development costs for similar functionality; or developers may seek to build complementary products that integrate with Apple's offering. While the present analysis cannot distinguish among these channels, the net effect is that Apple's entry shapes the direction of subsequent development activity, concentrating developer effort around capabilities Apple has signaled as priorities.

3.2 Incumbent Responses

I next examine how incumbent apps respond to Apple's entry across monetization and quality dimensions. All analyses in this subsection use a balanced monthly panel of apps operating on the platforms for the 9 months before and after Apple's entry. My primary estimator for incumbent app analyses is again the interaction-weighted estimator of Sun and Abraham (2021), which aggregates cohort-specific treatment effects while directly estimating pre-treatment event-time coefficients. This estimator addresses potential bias from heterogeneous treatment effects and staggered adoption by using never-treated and not-yet-treated units as controls, avoiding the negative-weighting problems that can afflict standard two-way fixed effects (TWFE) when already-treated units serve as controls for newly-treated units. Standard errors are clustered at the app level. For comparison, I also report results from standard TWFE and the Extended TWFE (ETWFE) estimator of Wooldridge (2021) in Appendix A. Results are qualitatively similar across methods.

Notably, not all apps in an affected market are equally "treated" by platform entry. An app providing nearly identical functionality to Apple's entrant faces different competitive pressure than a tangentially related app at the market periphery. I exploit the continuous similarity measure from the SBERT embeddings to estimate how treatment effects vary with competitive proximity. I use B-splines to flexibly estimate this relationship.

For each app i in market a, I observe its scaled similarity $\tilde{d}_i \in [\theta, 1]$ to the entrant, where $\theta = 0.6$ is the market inclusion threshold. I model treatment intensity as a function of this similarity:

$$y_{it} = \sum_{k=1}^{K} \phi_k \Big(\text{Post}_{at} \times \text{iOS}_i \cdot B_k(\tilde{d}_i) \Big) + \gamma_i + \eta_t + \varepsilon_{it}, \tag{5}$$

where $B_k(\cdot)$ are B-spline basis functions and K is the number of basis functions (degrees of freedom). B-splines offer several advantages: (1) they provide smooth, continuous functions without imposing specific functional forms; (2) they nest linear and polynomial specifications as special cases; (3) they avoid the boundary problems of polynomial regression; and (4) they allow for non-monotonic relationships.

I select K via 10-fold cross-validation at the app level, testing values from 3 to 8 and choosing the specification that minimizes normalized root mean squared error. This data-driven approach

balances model flexibility against overfitting. From the estimated spline, I compute the conditional treatment effect at scaled similarity $\tilde{d}$ as

$$\tau(\tilde{d}) \equiv \frac{\partial\, \mathbb{E}[y_{it} \mid \text{Post}_{at} \times \text{iOS}_i,\, \tilde{d}_i = \tilde{d}]}{\partial\, \text{Post}_{at} \times \text{iOS}_i} = \sum_{k=1}^{K} \phi_k\, B_k(\tilde{d}). \tag{6}$$

I interpret $\tau(\tilde{d})$ as the treatment effect at a given similarity level to the entrant. For presentation, I evaluate and plot $\tau(\tilde{d})$ on a grid $\tilde{d} \in [\theta, 1]$ together with pointwise 95% confidence intervals (standard errors clustered at the app level). This profile shows how competitive pressure varies with similarity, making it possible to identify a "competitive radius" where effects become statistically indistinguishable from zero.

3.2.1 Monetization Responses

Apple's entry can affect incumbent apps' monetization strategies through three main channels: app prices (measured in dollars for all apps including those charging $0), the proportion of apps offered for free, and in-app purchase adoption (whether an app offers additional paid content within the app).

Table 6 presents estimates for the monetization outcomes. I find that price increases by $0.63 following entry, representing approximately a 70% increase relative to the pre-treatment mean of $0.87. This aggregate price increase masks important compositional changes. Apps become 8 percentage points less likely to be free—representing approximately a 10% decrease in free offerings. However, conditional on charging a positive price, the price effect is modest ($0.05), suggesting the overall price increase is driven primarily by apps switching from free to paid rather than existing paid apps raising prices.

Table 6: Incumbent App Outcomes

	Price	Price (Paid Apps)	Free App	In-App Purchases	Rating	Update	Rating Count (IHS)
ATT	0.6306***	0.0531	-0.0797***	-0.0203**	0.0127	-0.0017	-0.0371
	(0.0380)	(0.0398)	(0.0104)	(0.0085)	(0.0080)	(0.0132)	(0.0247)
Baseline Mean	0.8704	4.040	0.7846	0.3807	3.983	0.2501	6.630
N	85,482	25,585	85,482	67,986	81,594	85,482	85,482

Notes: Sun & Abraham (2021) interaction-weighted DiD estimates at the app-month level. For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors clustered at the app level are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

The shift toward paid models appears to be a defensive response, with previously free apps beginning to charge rather than paid apps increasing prices. In-app purchase (IAP) adoption declines by 2.0 percentage points under the SA estimator. One interpretation is that developers simplify monetization strategies under competitive pressure, moving from complex freemium models to straightforward paid downloads, or that they move monetization from on-platform IAPs to off-platform subscriptions, which are not subject to Apple's typical commission.

Figure 3 presents event studies for all monetization outcomes. Pre-trends are generally flat under the SA estimator. For some outcomes (particularly price and IAP), TWFE shows pre-trend patterns that are not present under SA, reflecting concerns about heterogeneous effects across staggered entry events. Post-treatment estimates are generally aligned across estimators, with SA producing estimates that fall between TWFE and ETWFE.

The overall price increase begins immediately at entry and continues growing through the end of the study window, reaching approximately \$0.70. The free proportion shows a a large initial decline, which then settles out at a less extreme, but still suppressed level. The conditional price effect for paid apps remains relatively stable, indicating the overall price increase is driven by the composition change. Finally, IAP adoption decreases sharply at entry and remains suppressed.

The immediate response suggests developers quickly recognize the competitive threat and adjust monetization strategies without extended experimentation, at least beyond the first month. The continued evolution through the post-period indicates either learning about optimal responses or cascading effects as multiple developers adjust strategies simultaneously.

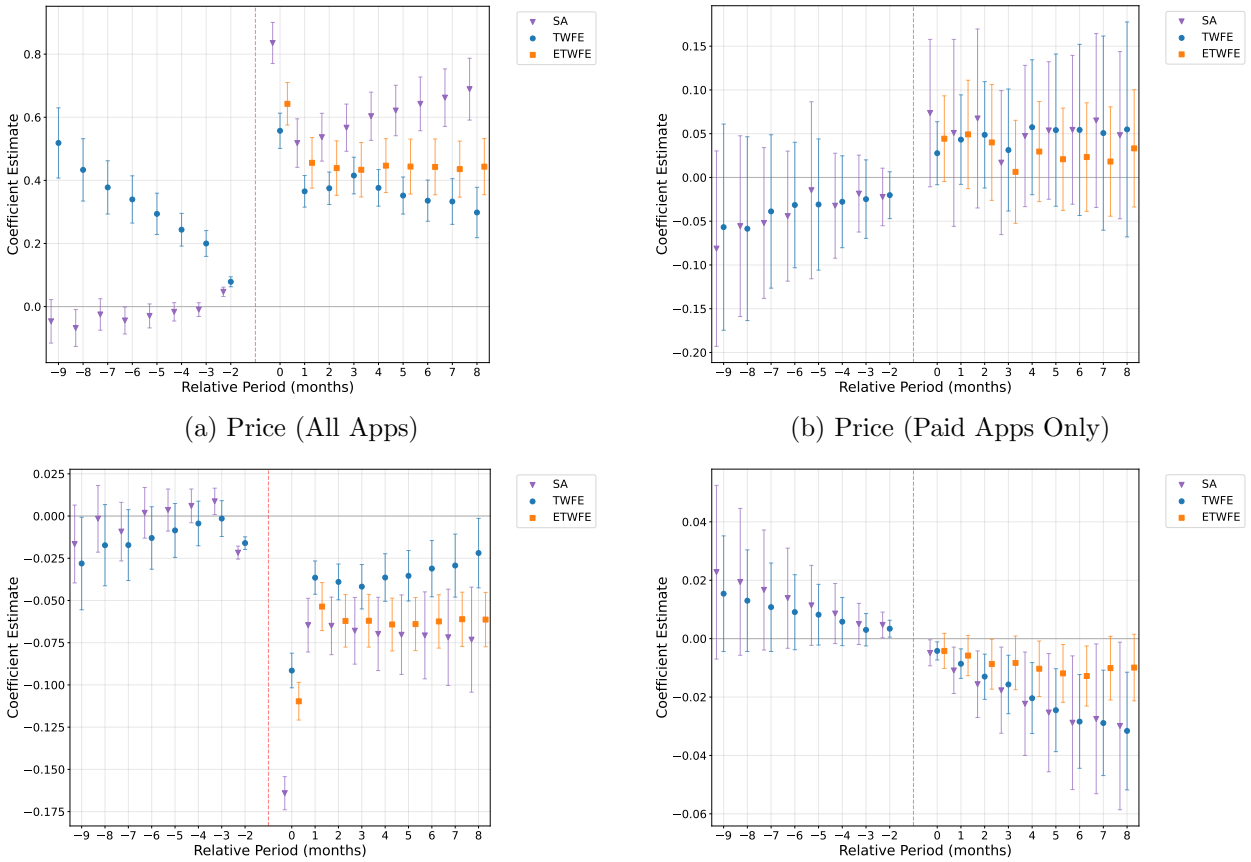
For market-specific treatment effects, I use TWFE with market-indicator interactions rather than SA.

$$y_{it} = \sum_{a \in \mathcal{A}} \alpha_a \cdot \text{Post}_{at} \times \text{iOS}_i + \gamma_i + \eta_t + \varepsilon_{it} \tag{7}$$

where α_a captures the market-specific treatment effect.

Because multiple markets can share the same treatment timing, the SA estimator's cohort-based aggregation does not distinguish heterogeneity across markets—only across treatment cohorts. Figure 4 reveals substantial heterogeneity in monetization responses across markets. Price effects vary substantially in magnitude, with many markets showing limited or no response. Free proportions decline in most markets, though magnitudes vary.

Figure 3: Monetization Dynamics



(a) Price (All Apps)

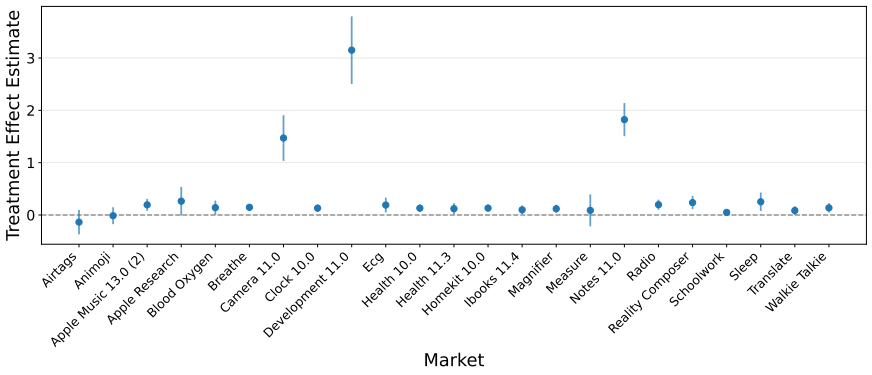
(b) Price (Paid Apps Only)

(c) Free

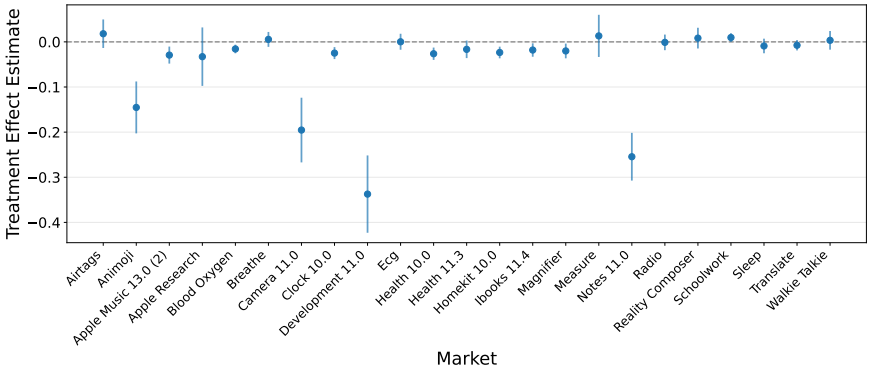
(d) In-App Purchases

Note: Event study estimates showing dynamic treatment effects relative to one month before entry. SA (Sun and Abraham (2021)) is the primary estimator; TWFE and ETWFE (Wooldridge (2021)) shown for robustness. Analysis is at the app-month level. For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors are clustered at the app level. Error bars represent 95% confidence intervals.

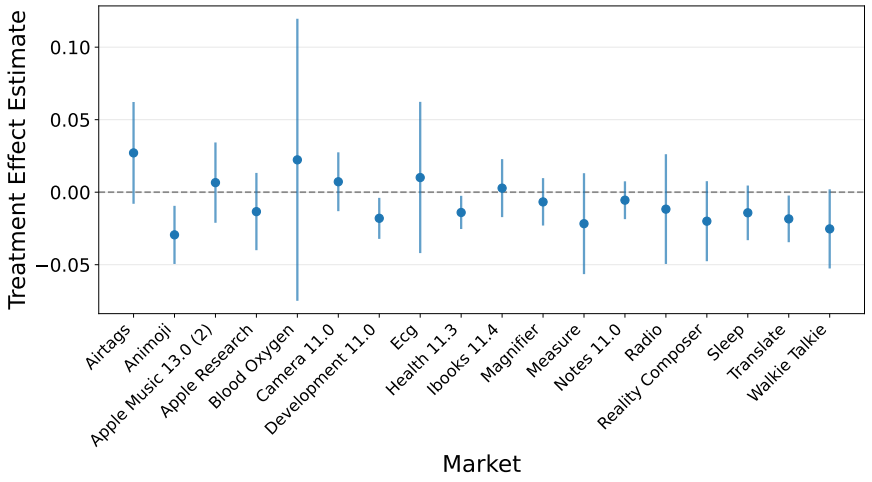
Figure 4: Market Heterogeneity in Monetization



(a) Price (All Apps)



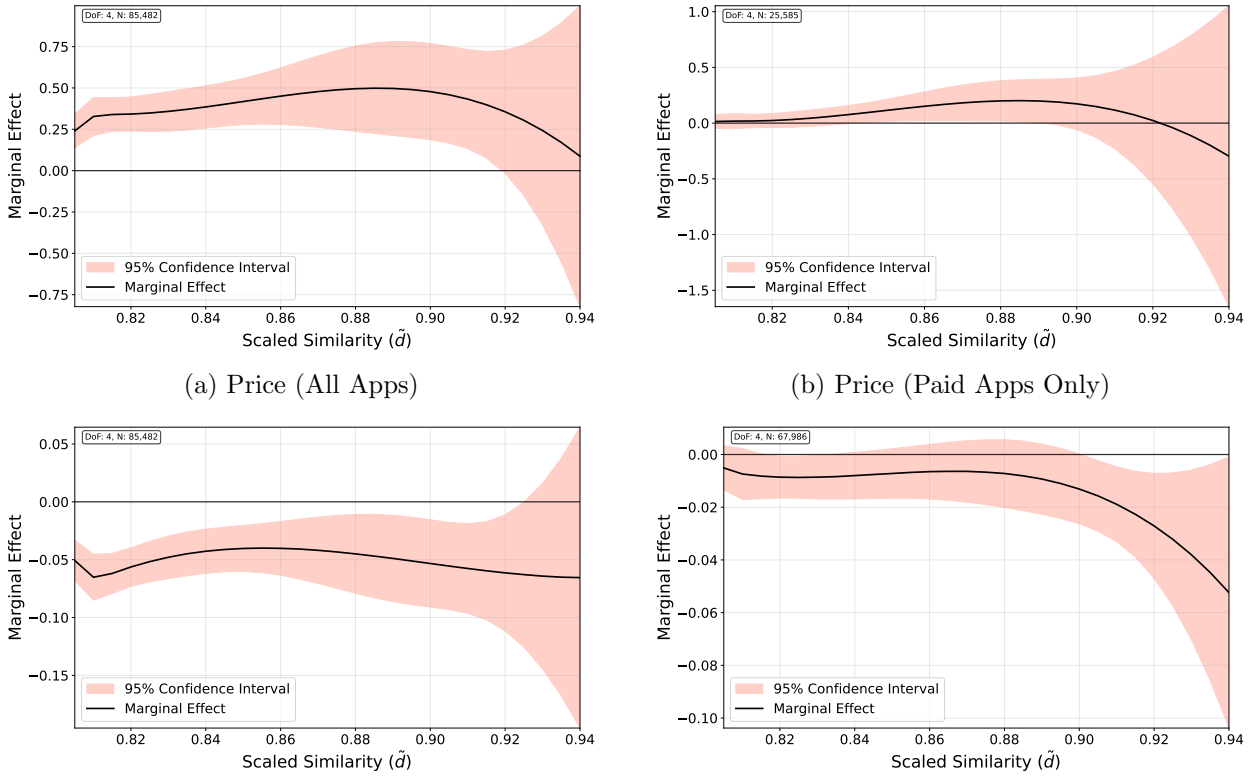
(b) Free



(c) In-App Purchases

Note: Market-specific treatment effects from TWFE model. Analysis is at the app-month level. For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors are clustered at the app level. Points show estimates with 95% confidence intervals for each of 23 markets.

Figure 5: Monetization Effects by Distance from Entrant



(a) Price (All Apps)

(b) Price (Paid Apps Only)

(c) Free

(d) In-App Purchases

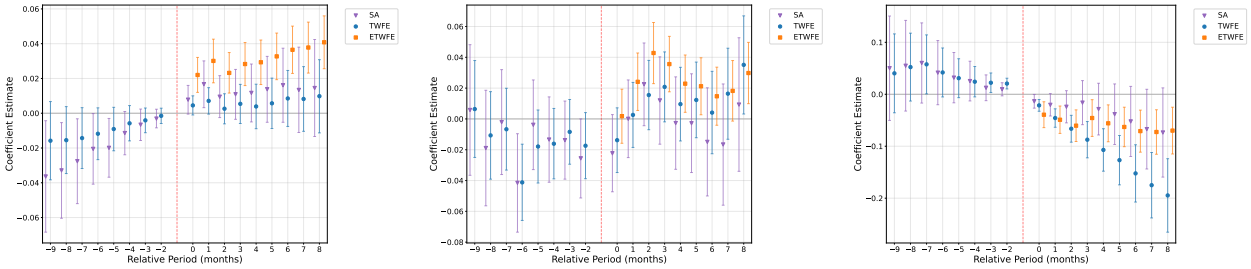
Note: Marginal treatment effects across competitive distances from B-spline distance regression (Equation (5)) with 4 degrees of freedom, where treatment intensity varies with competitive similarity to Apple's entrant. Analysis is at the app-month level. For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors are clustered at the app level. Shaded areas represent 95% confidence intervals. $\tilde{d} = 1$ represents entrant location.

Figure 5 presents marginal effects showing how monetization responses vary with competitive distance from the entrant, which is located at scaled similarity, $\tilde{d} = 1$. While the price effect is positive throughout the study domain, the effect is the largest in the middle region—among apps that are not most nor least exposed to Apple's entrant. The effect of entry on IAP adoption varies substantially with distance. Apps farthest from Apple's entrant appear unaffected, with IAP adoption consistently falling to a larger degree as distance decreases.

3.2.2 Quality Responses

Apple's entry can affect incumbent app quality through three dimensions: app ratings (average user star ratings on a 1-5 scale), update frequency (as a measure of development intensity and

Figure 6: Quality Dynamics



(a) Rating (b) Update Frequency (c) Rating Count (IHS)

Note: Event study estimates showing dynamic treatment effects relative to one month before entry. SA (Sun and Abraham (2021)) is the primary estimator; TWFE and ETWFE (Wooldridge (2021)) shown for robustness. Analysis is at the app-month level. Standard errors are clustered at the app level. Error bars represent 95% confidence intervals.

ongoing improvement), and rating count (IHS-transformed to serve as a proxy for app demand and user engagement).

Table 6 presents estimates for quality outcomes. Under the SA estimator, average ratings show a modest and imprecisely estimated improvement of 0.01 stars—economically small relative to a base of approximately 3.8 stars, and unlikely to meaningfully affect consumer choice.

Update frequency shows essentially no effect, with a point estimate near zero. This suggests that, on average, developers neither increase nor decrease development intensity in response to Apple's entry, though this average may mask heterogeneous responses across apps.

Rating count (IHS transformed), a proxy for demand, shows a negative effect of -3.7% though also not statistically significant. The pattern—stable or slightly improved quality but potentially lower demand—suggests that increased competition from Apple may reduce visibility or market share for incumbents even as they maintain their offerings. This finding stands in contrast to the positive effect on entering cohorts, suggesting that incumbents may not be repositioning their products enough in response to Apple's entry.

Figure 6 presents event studies for all quality outcomes. Rating improvements emerge gradually and continue to grow through the end of the study window. Update frequency shows minimal changes initially but exhibit a small, and seemingly transitory, rise roughly four months post-entry. Rating count shows a persistent decline post-entry.

Figure 7 presents market-level differences in the quality outcomes. The estimated effects of

entry on ratings are positive, but imprecisely estimated, for almost all markets. Similarly for the effect on product updating. There is greater variation, both in sign and magnitude in the effect on rating count, which proxies consumer demand and product use intensity. In a few cases we observe expansionary effects, but the average effect is generally close to zero or negative.

Figure 8 presents distance gradients for quality outcomes. Quality responses show weaker distance gradients than monetization outcomes. Apps closest to the entrant exhibit sharp declines in ratings, with the negative effect diminishing with distance. Eventually, the sign flips, as more peripheral products see ratings improvements. This suggests quality competition varies dramatically with distance, possibly reflecting the fact that the quality (or perceived quality) of the closest competitors is harmed by first-party entry, while more distance apps may benefit from complementaries with the entrant or related improvements in the overall software and hardware platform. Update frequency follows a similar pattern. Rating count shows no clear pattern.

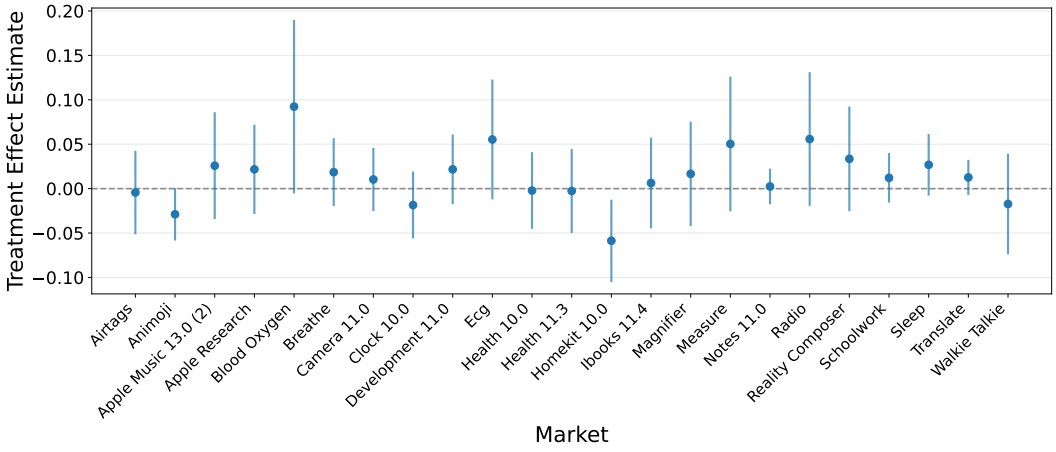
3.2.3 Discussion: Two Key Dimensions of Heterogeneity

Section 3.1–Section 3.2.2 document substantial heterogeneity along two independent dimensions. First, effects vary *across markets*: some markets show positive responses (expansion, quality improvements) while others show negative responses (contraction, defensive pricing). Notably, many show no response. the market heterogeneity figures reveal that for most outcomes, effects range from strongly negative to strongly positive across the 23 markets, with both signs well-represented.

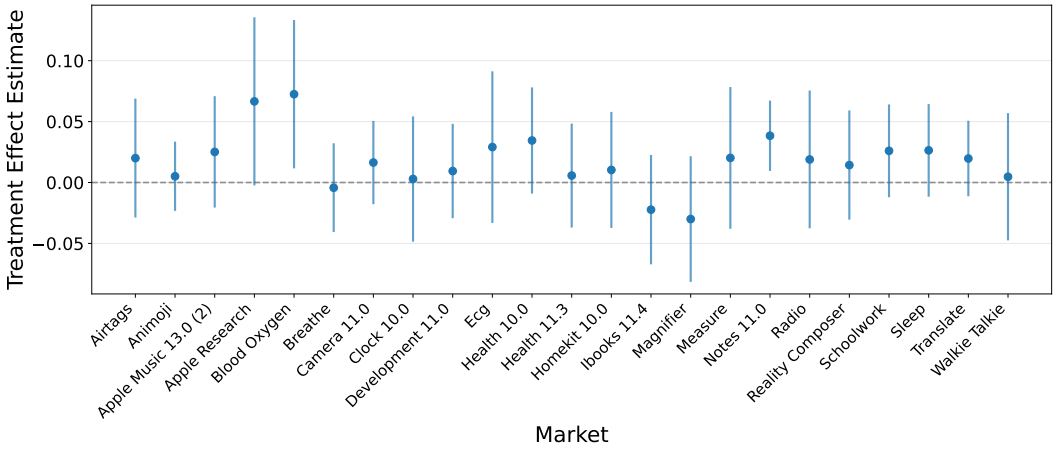
Table 7 quantifies this heterogeneity, showing that across all outcomes, 66% of markets experience no statistically significant effects from Apple's entry, indicating that many markets are largely unharmed by first-party entry. Monetization outcomes show the most consistent responses: 87% of markets experience price increases and 52% see declines in free app proportions. In contrast, quality outcomes are largely unaffected, with 87% of markets showing no significant change in ratings and update frequency. These patterns reveal that platform entry's competitive effects are concentrated in pricing and business model decisions rather than quality competition.

Second, effects vary *within markets* by competitive proximity. Apps closer to the entrant in product space experience larger treatment effects than distant apps. This creates a "competitive radius" around each entrant, beyond which effects dissipate. Importantly, distance gradients themselves vary across markets—some show steep decay (proximity matters greatly) while others show

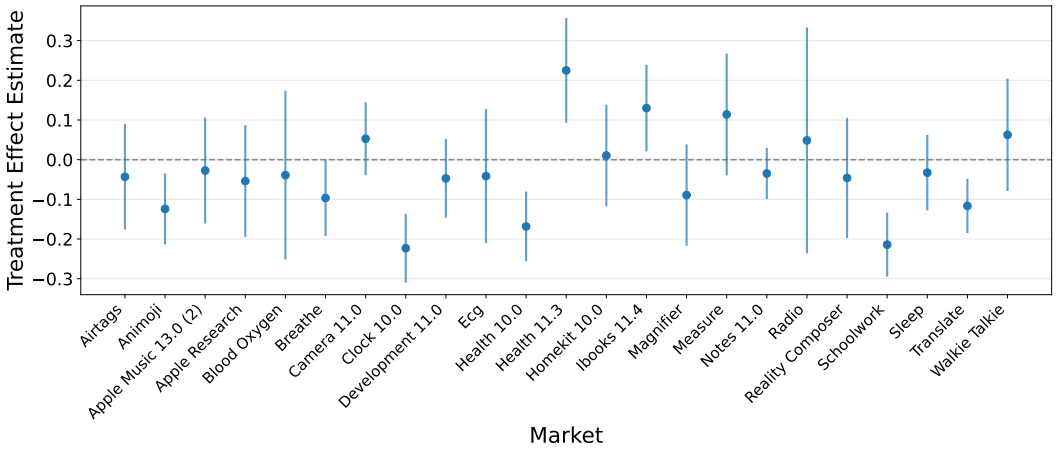
Figure 7: Market Heterogeneity in Quality



(a) Rating



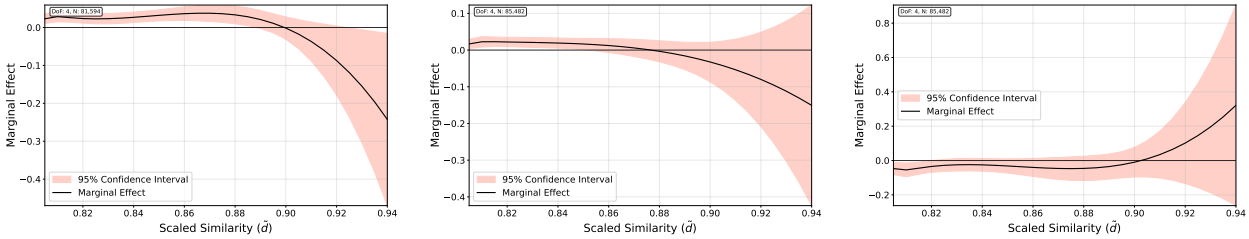
(b) Update Frequency



(c) Rating Count (IHS)

Note: Market-specific treatment effects from TWFE model. Analysis is at the app-month level. Standard errors are clustered at the app level. Points show estimates with 95% confidence intervals for each of 23 markets.

Figure 8: Quality Effects by Distance from Entrant



(a) Rating (b) Update Frequency (c) Rating Count (IHS)

Note: Marginal treatment effects across competitive distances from B-spline distance regression (Equation (5)) with 4 degrees of freedom, where treatment intensity varies with competitive similarity to Apple's entrant. Analysis is at the app-month level. Standard errors are clustered at the app level. Shaded areas represent 95% confidence intervals. $\tilde{d} = 1$ represents entrant location.

Table 7: Distribution of Market-Level Treatment Effects

Outcome	Percentage of Markets			Dominant Pattern
	Positive	Negative	Null	
Monetization				
Price	87%	0%	13%	Positive
Free	4%	52%	43%	Negative
In-App Purchases	0%	22%	78%	No effect
Price (Paid Apps)	13%	0%	87%	No effect
Quality				
Rating	4%	9%	87%	No effect
Update Frequency	13%	0%	87%	No effect
Rating Count (IHS)	9%	26%	65%	No effect
Average	**19%**	**16%**	**66%**	–

Notes: Based on market-level TWFE estimates for 23 markets. Effects are classified as positive or negative if statistically significant at the 10% level. "Null" indicates no statistically significant effect. Dominant pattern shows the most common effect type for each outcome.

relatively flat gradients (proximity matters less).

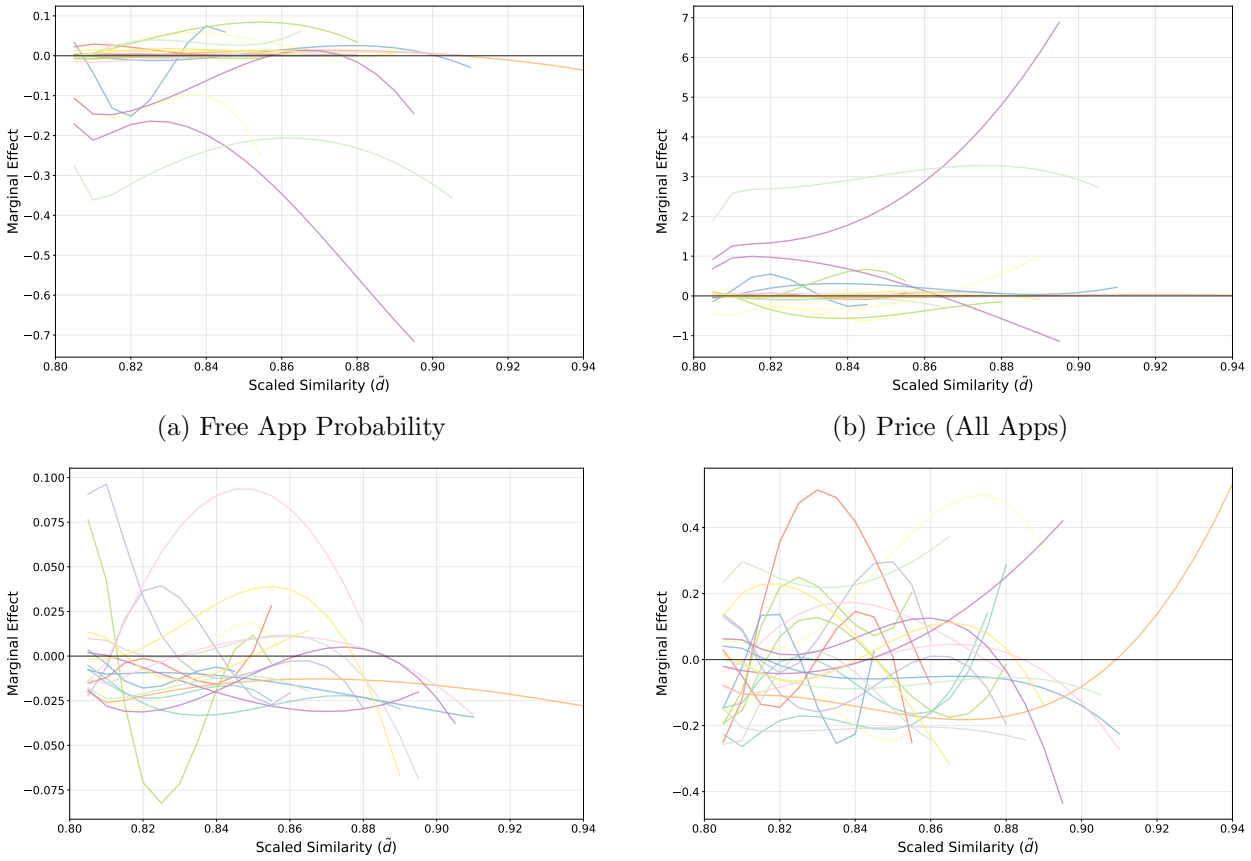
Figure 9 reveals the full extent of heterogeneity by showing distance gradients for all 23 markets simultaneously. Because markets differ in their distribution of competitive distances, each market's curve is displayed only within its observed similarity range, avoiding extrapolation beyond the data support. The free outcome panel shows remarkable variation among affected markets: some exhibit steep declines where apps become substantially less likely to be free at close distances but show no effect at far distances, while others maintain relatively constant effects across all distances. Price effects display a similar pattern of heterogeneity.

The IAP and rating count panels reveal contrasting patterns. IAP adoption declines are relatively consistent across distances within markets that show effects, though the magnitude varies dramatically across markets. Rating count effects are more dispersed, with some markets showing increased engagement for close competitors while others show uniform declines across all distances. In both cases, markets also vary in the sign of the estimated effect. This variation in competitive patterns across outcomes suggests that platform entry affects different aspects of app strategy through distinct mechanisms.

Several striking patterns emerge from this comprehensive view of heterogeneity. First, as documented in Table 7, a substantial proportion of markets show essentially no response to Apple's entry. This suggests that platform entry often has limited competitive impact, contrary to concerns about systematic foreclosure. Second, among markets with more dramatic effects, monetization outcomes (free, price, IAP) show more consistent directional patterns, while quality outcomes (rating count) show bidirectional responses. Third, the "competitive radius" varies dramatically across affected markets: some show effects only for particularly close apps, while others affect a much broader set of apps.

The contrast between outcomes is particularly revealing. The free and price panels show that monetization responses are often complementary—markets with steep free gradients tend to show steep price gradients, suggesting coordinated strategic shifts from free to paid models among close competitors. The IAP panel shows more uniform within-market effects but dramatic cross-market variation in magnitude, consistent with market-specific factors (e.g., user demographics, app categories) determining the viability of freemium models. The rating count dispersion suggests that demand responses depend on factors beyond just competitive distance, possibly including network

Figure 9: Distance Gradients Across All Markets



(a) Free App Probability

(b) Price (All Apps)

(c) In-App Purchases

(d) Rating Count (IHS)

Note: Market-specific marginal effects across competitive distances from B-spline distance regression (Equation (5)) with 4 degrees of freedom. Analysis is at the app-month level. For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors are clustered at the app level. Each line represents one market's treatment effects across distance, with $\tilde{d} = 1$ representing entrant location. Curves are restricted to each market's observed similarity range to avoid extrapolation beyond the data.

effects, switching costs, or Apple's market expansion effects.

These patterns have important implications for policy. The prevalence of null effects challenges narratives of universal competitive harm from platform entry. The concentration of effects among close competitors (when effects exist at all) suggests that most apps in a market remain unaffected, limiting the scope of potential welfare losses. The divergence across outcomes indicates that platform entry triggers complex strategic adjustments rather than simple foreclosure. Heterogeneity across markets implies that blanket regulations would be both over- and under-inclusive—restricting beneficial entry in some markets while failing to address harm in others. The two-dimensional nature of heterogeneity argues for nuanced, market-specific regulatory approaches that consider both the type of entry and the competitive structure of affected markets.

4 Mechanisms and Interpretation

I have documented substantial heterogeneity across markets and within markets by distance. What explains the cross-market variation in average effects? I show that entry type—whether Apple enters via a new app (standalone) or OS integration (integrated)—accounts for meaningful variation in observed responses. Note that distance analysis (Section 3.2.3) already explains within-market variation; here I focus on across-market differences.

4.1 Entry Type Matters

Table 8 compares average treatment effects by entry type. The 14 standalone entries (new apps) and 9 integrated entries (OS integrations) generate systematically different competitive responses across nearly all outcomes.

Two key patterns emerge from Table 8. First, integrated entry generates dramatically larger monetization effects. Price increases are nearly 16-fold larger for integrated ($1.26) than standalone ($0.08) entry. The proportion of free apps declines by 17.8 percentage points for integrated versus only 2.6 percentage points for standalone—a seven-fold difference. These dramatic differences suggest fundamentally different competitive dynamics: integrated entry forces aggressive monetization responses, while standalone entry generates modest adjustments.

Second, both entry types reduce demand, but integrated entry causes much larger declines.

Table 8: [Preliminary] Entry Type Heterogeneity: Standalone vs Integrated

Outcome	All	Standalone	Integrated	Difference
Price	0.379***	0.076***	0.179	-0.103
	(0.0277)	(0.0197)	(0.2116)	(0.2125)
Free	-0.040***	-0.024***	-0.052	0.028
	(0.0067)	(0.0075)	(0.0332)	(0.0340)
In-App Purchases	-0.019***	-0.029**	-0.010	-0.018
	(0.0059)	(0.0114)	(0.0577)	(0.0588)
Price (Paid Apps)	0.047	0.061	0.186	-0.124
	(0.0379)	(0.0413)	(0.6432)	(0.6445)
Update	0.011	0.0014	-0.0046	0.0060
	(0.0102)	(0.0148)	(0.0676)	(0.0692)
Rating	0.0062	-0.0036	-0.025	0.021
	(0.0060)	(0.0091)	(0.0631)	(0.0638)
Rating Count (IHS)	-0.108***	-0.105***	-0.412***	0.307*
	(0.0202)	(0.0275)	(0.1584)	(0.1608)

Notes: TWFE estimates with standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. All: 23 markets (14 standalone + 9 integrated). Difference: Standalone - Integrated. Standard errors for differences calculated assuming independence.

This contradicts simple market expansion theories and suggests that the unavoidable nature of integrated features creates particularly severe competitive pressure.

One possible explanation for these differences is the unavoidability of integrated releases. OS-integrated features cannot be separately uninstalled or avoided by users, creating fundamentally different competitive pressure than marketplace apps. When Apple releases a standalone app like Translate, users must actively discover it in the App Store, evaluate it against alternatives, and choose to download it. Market forces operate normally—if third-party alternatives are superior, users may never adopt Apple's offering.

Integrated entry eliminates this choice architecture. When Apple integrates translation capabilities directly into iOS, every user automatically possesses this functionality upon updating their device. Third-party translation apps must now compete against a feature that users already have, cannot remove, and may discover through system prompts or Siri suggestions. This unavoidability likely explains the larger monetization responses to integrated entry: facing an unremovable competitor, third-party apps must differentiate more aggressively through pricing and business model changes.

The unavoidability mechanism also explains the immediate response timing observed in event studies. While standalone entry effects often emerge gradually as users discover and adopt Apple's app, integrated entry effects appear immediately when the OS update deploys to millions of devices near-simultaneously.

The cohort analysis reveals that entry type differences extend beyond incumbent responses to affect the composition of new entrants. Comparing cohort ATT estimates for standalone markets versus integrated markets shows that the patterns documented above—larger effects for integrated entry—also appear in entry characteristics. New apps entering integrated markets after Apple's entry show more pronounced shifts in monetization strategies and quality characteristics than those entering standalone markets. This suggests that anticipated competitive pressure shapes not just incumbent behavior but also entry decisions, with developers either deterred from entering or strategically repositioning when entering markets where Apple has integrated functionality. The consistency of entry type effects across both incumbent responses and cohort composition strengthens the interpretation that unavoidability fundamentally changes competitive dynamics.

5 Concluding Discussion

This paper provides systematic evidence on how platform-owner entry affects third-party developers across multiple markets. Studying 23 instances of Apple entering submarkets in its App Store—14 through standalone app releases and 9 through integrated OS features—I uncover striking heterogeneity that challenges both pro-platform and anti-platform narratives. The evidence reveals that platform entry is neither uniformly beneficial nor uniformly harmful, but rather generates widely varying effects that depend systematically on market context, entry implementation, and competitive proximity.

Perhaps most surprisingly, I find that many of markets experience no statistically significant effects on average. Across all outcomes, 66% of market-level estimates are statistically indistinguishable from zero, with quality metrics particularly unaffected—87% of markets show no significant changes in ratings and update frequency. This prevalence of null effects contradicts popular claims of either universal harm from foreclosure or universal benefits from technology spillovers. When effects do occur, they largely concentrate in monetization strategies, suggesting that platform entry primarily triggers business model adjustments rather than changes to product development.

The heterogeneity I document operates along three key dimensions. First, effects vary meaningfully across markets, with some experiencing increased entry and quality improvements while others face reduced entry and defensive pricing. Second, within markets, effects vary with the similarity to Apple's entrant. Third, entry type matters fundamentally: integrated OS features generate effects that are an order of magnitude larger than standalone apps. Several important limitations qualify these findings and their policy implications. Most fundamentally, the reduced-form approach estimates net competitive effects but cannot directly assess consumer welfare. The combination of and variation in effects documented here have ambiguous welfare implications. For example, apps changing from free to paid may also reduce reliance on intrusive advertising or data collection.

The measurement challenges inherent in platform markets further constrain this analysis. I observe only nine months of post-entry effects due to balanced panel requirements, potentially missing longer-run dynamics where initial defensive responses evolve into accommodation, exit, or innovation. The shift toward paid apps that I document may involve important but unobserved changes

in advertising intensity, data collection practices, or on- or off-platform subscription models that affect consumer welfare in ways this analysis cannot capture. Rating counts provide an imperfect proxy for demand.

Despite these limitations, these findings open important avenues for future research. The stark heterogeneity I document—both across markets and within markets by competitive distance—calls for deeper investigation into what determines these differential effects. Understanding why 85% of market-outcome pairs show positive or null effects while others experience negative changes could help predict which markets are most vulnerable to platform competition and inform more targeted regulatory approaches. The systematic differences between standalone and integrated entry also warrant further study, particularly as platforms increasingly blur the lines between operating system features and marketplace applications.

The longer-run dynamics of platform competition remain largely unexplored. The nine-month post-entry window may miss important evolutionary patterns where initial defensive responses give way to accommodation, innovation, or exit. Technology spillovers from platform entry might take years to fully materialize, particularly when new APIs and frameworks require developers to learn new skills and rebuild applications. Understanding these dynamic patterns could reveal whether the immediate competitive harm from platform entry is offset by longer-run innovation benefits, or whether initial advantages compound over time into durable market dominance.

These findings provide crucial empirical evidence for the ongoing global debate over platform regulation and antitrust enforcement. The evidence reveals that calls for blanket restrictions (see, for example, Warren (2019)) fundamentally misunderstand the nature of platform competition. The prevalence of null effects demonstrates that categorical bans would prevent substantial platform entry that causes limited to no competitive harm whatsoever. More troublingly, the concentration of effects among close competitors means that broad statutory language could restrict entry that affects only a narrow segment of apps while leaving the vast majority unaffected.

The stark differences between standalone and integrated entry provide actionable guidance for antitrust enforcement and regulatory design. Integrated OS features present fundamentally different competitive concerns than standalone apps that users can choose to ignore. This distinction maps directly onto antitrust theory: integrated features resemble tying arrangements or technological bundling that forecloses competition, while standalone apps are more likely to compete on merit

in the marketplace.[7] Antitrust authorities should therefore apply different levels of scrutiny based on entry mode, with integrated features triggering heightened review similar to merger analysis in adjacent markets. The unavoidability of integrated features creates the kind of foreclosure concerns that antitrust law is designed to address, while standalone apps' market-based adoption provides the competitive constraints that normally obviate regulatory intervention. Furthermore, the competitive radii documented in this paper have implications for regulatory and antitrust market definitions. The methods employed here provide an empirically grounded approach to defining relevant markets in platform cases—an often difficult challenge.

The heterogeneity across markets further challenges the premise underlying many current regulatory proposals—that platform entry is categorically harmful to innovation and competition. Rather than categorical entry prohibitions, policy should focus on conduct remedies that address specific mechanisms of harm: prohibiting self-preferencing in search rankings, mandating equal API access, requiring transparency in data usage, and ensuring user choice for default applications. These targeted interventions could prevent the severe demand declines I document for integrated features while preserving the innovation spillovers that benefit some markets. The evidence strongly supports a shift from the blunt instrument of entry bans to the precision tools of conduct regulation, guided by empirical evidence on which types of entry in which contexts generate competitive harm.

Ultimately, as policymakers worldwide grapple with platform power, evidence-based approaches that recognize the fundamental heterogeneity in platform competition effects will be essential for crafting interventions that preserve innovation incentives while preventing competitive harm. The complex, multi-dimensional nature of platform competition revealed in this study underscores that simple narratives—whether platforms are inherently good or bad for competition—fail to capture the nuanced reality of digital markets.

References

Anderson, Simon and Özlem Bedre Defolie. 2024. "Hybrid platform model: monopolistic competition and a dominant firm." *The RAND Journal of Economics* 55 (4):684–718. URL

[7] Although, pre-installed standalone apps may warrant greater scrutiny.

https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12478. 2

Boudreau, Kevin J. and Andrei Hagiu. 2011. "Platform Rules: Multi-Sided Platforms as Regulators." In *Platforms, Markets, and Innovation*, edited by Annabelle Gawer. Edward Elgar Publishing, Inc. 5

Congress.gov. 2022. "Text – S. 2992 – 117th Congress (2021–2022): American Innovation and Choice Online Act." https://www.congress.gov/bill/117th-congress/senate-bill/2992/text. Accessed 23 June 2025. 2

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer. 2019. "Competition policy for the digital era." URL https://data.europa.eu/doi/10.2763/407537. 2

De Chant, Tim. 2022. "Antitrust bill that bars Big Tech self-preferencing advances in Senate." *Ars Technica* URL https://arstechnica.com/tech-policy/2022/01/antitrust-bill-that-bars-big-tech-self-preferencing-advances-in-senate/. 2

Ershov, Daniel. 2024. "Variety-Based Congestion in Online Markets: Evidence from Mobile Apps." *American Economic Journal: Microeconomics* 16 (2):180–203. URL https://www.aeaweb.org/articles?id=10.1257/mic.20200347. 5

EUR-Lex. 2022. "Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act)." http://data.europa.eu/eli/reg/2022/1925/oj. Official Journal L 265/1, 12 October 2022. Accessed 23 June 2025. 2

Foerderer, Jens, Thomas Kude, Sunil Mithas, and Armin Heinzl. 2018. "Does Platform Owner's Entry Crowd Out Innovation? Evidence from Google Photos." *Information Systems Research* 29 (2):444–460. 5

Hagiu, Andrei, Tat-How Teh, and Julian Wright. 2022. "Should platforms be allowed to sell on their own marketplaces?" *The RAND Journal of Economics* 53 (2):297–327. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12408. 2, 5

Jabr, Wael, Dominik Gutt, Jurgen Neumann, and Dennis Kundisch. 2020. "The Fate of the App: Economic Implications of Updating under Reputation Resetting." *Working Paper* . 5

Johnson, Garrett, Tesary Lin, James C. Cooper, and Liang Zhong. 2025. "COPPAcalypse? The YouTube Settlement's Impact on Kids Content." Working paper, SSRN. Available at SSRN: https://ssrn.com/abstract=4430334 or http://dx.doi.org/10.2139/ssrn.4430334. 5

Kesler, Reinhold. 2022. "The Impact of Apple's App Tracking Transparency on App Monetization." *Working Paper* . 5

Lam, H. Tai. 2023. "Platform Search Design and Market Power." Working Paper. 5

Lee, Kwok-Hao and Leon Musolff. 2023. "Two-Sided Markets Shaped By Platform-Guided Search." Working Paper. 5

Leyden, Benjamin T. 2023. "There's an App for That: Understanding Product Updating Under Digitization." *Working Paper* . 5

———. 2025. "Platform design and innovation incentives: Evidence from the product rating system on Apple's App Store." *International Journal of Industrial Organization* 99:103133. URL https://www.sciencedirect.com/science/article/pii/S0167718724000882. 5

Li, Ding and Hsin-Tien Tsai. 2022. "Mobile Apps and Targeted Advertising: Competitive Effects of Data Exchange." *Working Paper* . 5

Madsen, Erik and Nikhil Vellodi. 2025. "Insider Imitation." *Journal of Political Economy* 133 (2):652–709. URL https://doi.org/10.1086/732888. 5

Raval, Devesh. 2023. "Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box." Working Paper. 5

Reimers, Nils and Iryna Gurevych. 2019. "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks." In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. URL https://arxiv.org/abs/1908.10084. 6

Scott Morton, Fiona, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, A Douglas Melamed, and Jamie Morgenstern. 2019. "Market Structure and Antitrust Subcommittee Report." In *Stigler Committee on Digital Platforms: Final Report*,

edited by Luigi Zingales, Guy Rolnik, and Filippo Maria Lancieri. George J. Stigler Center for the Study of the Economy and the State, The University of Chicago Booth School of Business, 23–138. URL https://research.chicagobooth.edu/-/media/research/stigler/pdfs/digital-platforms---committee-report---stigler-center.pdf. 2

Sun, Liyang and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics* 225 (2):175–199. URL https://www.sciencedirect.com/science/article/pii/S030440762030378X. Themed Issue: Treatment Effect 1. 13, 15, 17, 20, 23, 38, 39, 40

Teng, Xuan. 2023. "Self-preferencing, Quality Provision, and Welfare in Mobile Application Markets." Working Paper. URL https://ssrn.com/abstract=4259639. 5

Warren, Elizabeth. 2019. "Here's how we can break up Big Tech." *Medium* URL https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c. 2, 33

Wen, Wen and Feng Zhu. 2019. "Threat of platform-owner entry and complementor responses: Evidence from the mobile app market." *Strategic Management Journal* 40 (9):1336–1367. 4

Wooldridge, Jeffrey M. 2021. "Two-Way Fixed Effects, the Two-Way Mundlak Regression, and Difference-in-Differences Estimators." Tech. rep. URL https://ssrn.com/abstract=3906345. SSRN working paper, version August 16, 2021. 15, 17, 20, 23, 38, 39, 40

Zhu, Feng and Qihong Liu. 2018. "Competing with complementors: An empirical look at Amazon.com." *Strategic Management Journal* 39 (10):2618–2642. 4

Appendices

A Comparison Across DiD Estimators

This appendix compares results across multiple difference-in-differences estimators to assess robustness to recent econometric concerns about two-way fixed effects (TWFE) with staggered treatment timing and heterogeneous treatment effects. The main text uses the Sun and Abraham (2021) interaction-weighted estimator as the primary specification. This appendix presents TWFE and extended TWFE (ETWFE) results for comparison.

A.1 Alternative Estimating Equations

Two-Way Fixed Effects (TWFE). The standard TWFE specification is:

$$y_{it} = \alpha \cdot \text{Post}_{at} \times \text{iOS}_i + \gamma_i + \eta_t + \varepsilon_{it} \tag{8}$$

where y_{it} is the outcome for unit i in month t, Post_{at} indicates periods after market a's treatment, iOS_i indicates the iOS platform, γ_i are unit fixed effects, and η_t are time fixed effects. Recent econometric literature has identified potential biases in TWFE estimation with staggered treatment timing when treatment effects are heterogeneous across cohorts or over time.

Extended Two-Way Fixed Effects (ETWFE). The extended TWFE estimator from Wooldridge (2021) saturates the model with treatment indicators and their interactions with covariates, providing consistent estimates under heterogeneous treatment effects. ETWFE models treatment effects as potentially varying by cohort and time since treatment, avoiding the negative weighting problems that can afflict standard TWFE. The estimator recovers cohort-time specific treatment effects, which can then be aggregated to form average treatment effects.

Sun & Abraham (SA). As described in the main text, the SA estimator uses never-treated and not-yet-treated units as controls, explicitly avoiding the use of already-treated units that can introduce bias through negative weights. The main text presents SA as the primary estimator; this appendix shows that results are qualitatively robust to using TWFE or ETWFE.

Table 9: DiD Point Estimates

Outcome	SA	TWFE	ETWFE
Free App	-0.0797***	-0.0403***	-0.0667***
	(0.0104)	(0.0067)	(0.0075)
In-App Purchases	-0.0203**	-0.0195***	-0.0091**
	(0.0085)	(0.0059)	(0.0044)
Rating Count (IHS)	-0.0371	-0.1085***	-0.0585***
	(0.0247)	(0.0202)	(0.0174)
Price	0.6306***	0.3786***	0.4648***
	(0.0380)	(0.0277)	(0.0419)
Price (Paid Apps)	0.0531	0.0469	0.0295
	(0.0398)	(0.0379)	(0.0263)
Rating	0.0127	0.0062	0.0312***
	(0.0080)	(0.0060)	(0.0059)
Update	-0.0017	0.0114	0.0235***
	(0.0132)	(0.0102)	(0.0062)

Notes: Analysis is at the app-month level. Columns show treatment effect estimates from TWFE (two-way fixed effects), ETWFE (extended two-way fixed effects; Wooldridge (2021)), and SA (interaction-weighted estimator; Sun and Abraham (2021)). For the IAP outcome, 5 markets are excluded due to insufficient Android in-app purchase data in the pre-period. Standard errors clustered at the app level are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

A.2 Incumbent Analysis Robustness

Table 9 presents estimates across all three estimators for incumbent app outcomes. Results are qualitatively consistent across methods.

A.3 Entry/Exit Analysis Robustness

Table 10 presents market-level entry and exit results across estimators. Entry effects are imprecisely estimated and vary in sign across methods, consistent with the main text's finding of no robust effect on third-party entry. Exit effects are consistently negative across all estimators.

Table 10: DiD Point Estimates

Outcome	SA	TWFE	ETWFE
Entry Count (IHS)	0.0452	0.0841	-0.1065
	(0.1091)	(0.1439)	(0.0674)
Exit Count (IHS)	-0.1820	-0.1510	-0.0850
	(0.1232)	(0.1363)	(0.0675)

Notes: Analysis is at the platform-market-month level. Columns show treatment effect estimates from TWFE (two-way fixed effects), ETWFE (extended two-way fixed effects; Wooldridge (2021)), and SA (interaction-weighted estimator; Sun and Abraham (2021)). Standard errors clustered at the platform-market level are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

B Cohort Analysis Event Studies

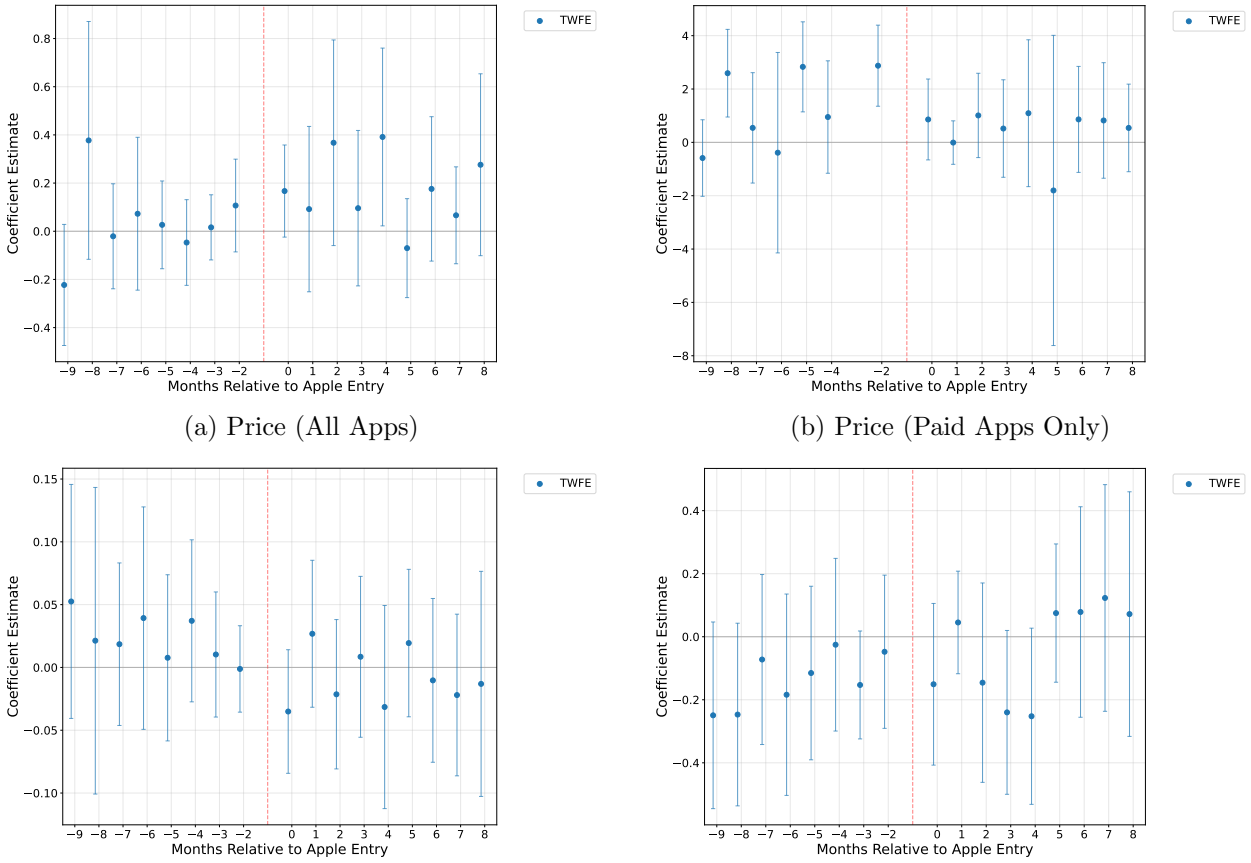
This appendix presents event study estimates for the cohort composition analysis described in Section 3.1.1. I estimate:

$$y_i = \sum_{k=-9}^{8} \beta_k \cdot \mathbf{1}[t_i - t_a^* = k] \times \text{iOS}_i + \delta_a + \eta_t + \varepsilon_i \tag{9}$$

where t_i is the month app i enters and t_a^* is Apple's entry month in market a. Coefficients β_k show how the iOS-Android gap in entry characteristics evolves relative to Apple's entry timing. This specification reveals whether changes in cohort composition occur gradually or abruptly around platform-owner entry.

Figures 10 and 11 present event study estimates for entering cohort characteristics.

Figure 10: Cohort Event Studies: Monetization Outcomes



(a) Price (All Apps)

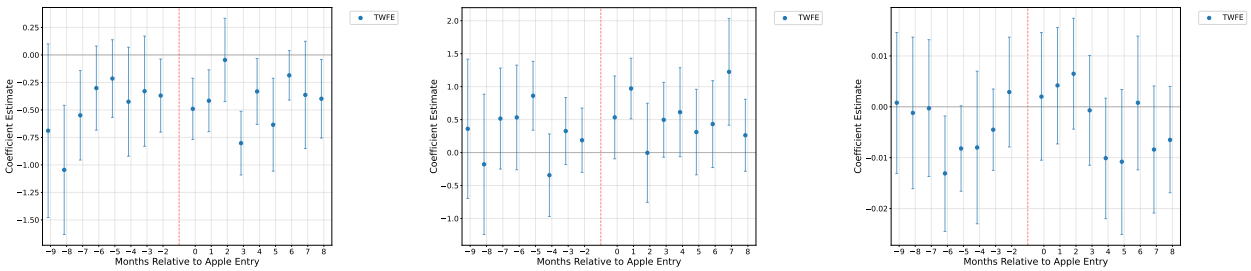
(b) Price (Paid Apps Only)

(c) Free

(d) In-App Purchases

Note: Event study estimates from cross-sectional DiD model showing how the iOS-Android gap in monetization characteristics evolves relative to Apple's entry timing. Each app is observed once at entry. Standard errors are clustered at the market-platform level. Error bars represent 95% confidence intervals.

Figure 11: Cohort Event Studies: Quality and Location Outcomes



(a) Rating

(b) Rating Count (IHS)

(c) Similarity

Note: Event study estimates from cross-sectional DiD model showing how the iOS-Android gap in quality and location characteristics evolves relative to Apple's entry timing. Each app is observed once at entry. Standard errors are clustered at the market-platform level. Error bars represent 95% confidence intervals.