Non-price Effects and Market Definition in Digital Markets

Junjun Zhang

Thesis presented for the degree of Doctor of Philosophy in Economics

School of Economics

University of East Anglia

06/2026

Abstract

This thesis investigates how mergers and acquisitions reshape outcomes of competition in digital markets, focusing on the non-price mechanisms. It combines basic theoretical modelling, structural estimation, and large-scale empirical causal analysis to study settings in which conventional price-based merger tools provide an incomplete account of competitive effects. Across three core chapters, the thesis shows that digital consolidation often operates through adjustments in perceived quality, innovation, and portfolio composition rather than through overt price movements. Since any assessment of merger and acquisition effects on competition depends on how the relevant market is defined, the thesis advances the literature by Natural Language Processing (NLP)-based approaches to digital market definition.

The first core chapter analyses Amazon's proposed acquisition of iRobot through a model of platform-driven quality allocation, in which a vertically integrated platform can selectively amplify perceived product quality via algorithmic rating and visibility contributions. Estimating a structural model from the United States (US) robotic vacuum cleaner market, it shows that even modest, algorithmically driven boosts to iRobot's ratings post-merger could substantially increase both the acquired firm and platform profits, reallocate market shares, and alter firms' incentives to invest in quality.

The second core chapter turns to the app economy and examines 1,387 Mergers and acquisitions (M&A) meticulously identified among US developers on the Google Play Store, using panel data on roughly 8 million apps and 1.8 million developers from 2015–2019. It documents that M&A reduce new app launches and external acquisitions while also limiting app removals and sales, leading to a net decline in product variety but an increase in maintenance updates and quality-enhancing activity on existing apps. The analysis further shows that acquisitions drive a sharper fall in new launches than mergers, and that horizontal and non-horizontal deals defined on product category exhibit distinct post-M&A strategies.

The third core chapter addresses two central questions in digital industrial organization: how to delineate relevant markets when user-side prices are often zero,

and how mergers within those markets reshape non-price conduct. Using Google Play data for 2015–2019, it develops an NLP-based framework to construct ‘text-defined markets’ from app descriptions and update notes, computes concentration measures in these markets, and then studies 1,914 horizontal M&A events, identified by the presence of apps from both the acquirer and the target in the same market, using a staggered Difference-in-Differences (DiD) design with matched controls. The findings show that horizontal mergers that would be flagged by European Union (EU) or US screening thresholds tend to increase developer exit (‘pruning’) and modestly raise innovation, without clear price effects, while below-threshold horizontal and non-horizontal integrations mainly reallocate update effort and slightly raise prices.

Contents

List of Acronyms

ΔHHI Change in the Herfindahl–Hirschman Index.

ARI Adjusted Rand Index.

ATT Average Treatment Effect on the Treated.

BLP Berry–Levinsohn–Pakes.

CMA Competition and Markets Authority.

CR4 four-firm concentration ratio.

CSDID Callaway and Sant'Anna staggered DiD estimator.

DBSCAN Density-Based Spatial Clustering of Applications with Noise.

DG COMP Directorate-General for Competition.

DiD Difference-in-Differences.

DMA Digital Markets Act.

DOJ Department of Justice.

EU European Union.

EUMR EU Merger Regulation.

FOC First Order Condition.

FTC Federal Trade Commission.

FTE Full-time equivalent.

FY Fiscal Year.

GAFAM Google, Apple, Facebook, Amazon, and Microsoft.

GDPR General Data Protection Regulation.

GMM Generalized Method of Moments.

GWB Gesetz gegen Wettbewerbsbeschränkungen.

HDBSCAN Hierarchical Density-Based Spatial Clustering of Applications with Noise.

HHI Herfindahl–Hirschman Index.

IV Instrumental Variables.

M&A Mergers and acquisitions.

NLP Natural Language Processing.

OECD Organisation for Economic Co-operation and Development.

OLS Ordinary Least Squares.

PSM Propensity Score Matching.

R&D Research and Development.

RCL Random Coefficients Logit.

SSNDQ Small but Significant Non-transitory Decrease in Quality.

SSNIP Small but Significant Non-transitory Increase in Price.

TWFE Two-Way Fixed Effects.

UK United Kingdom.

UMAP Uniform Manifold Approximation and Projection.

US United States.

USD United States Dollar.

List of Figures

List of Tables

Acknowledgements

I am deeply indebted to my supervisors, Michael Kummer and Franco Mariuzzo, for their unwavering guidance, intellectual generosity, and patience throughout the course of my doctoral studies. Their insight, encouragement, and high academic standards have shaped this thesis in more ways than I can fully express. This doctoral journey began in September 2021, at the height of the COVID-19 pandemic. Although I had received my offer and was preparing to commence my PhD, a combination of visa delays and international travel restrictions forced me to return to my home country. This period involved substantial financial cost and personal uncertainty, and it posed a serious challenge to the continuation of my studies. Throughout this time, my supervisors offered exceptional support and reassurance, and—together with the School of Economics at the University of East Anglia—worked to ensure that my scholarship could be retained and that my research could continue despite these disruptions. Their flexibility, understanding, and commitment were decisive in allowing me to remain academically active during this period. I formally joined UEA in February 2022, and I am profoundly grateful for the institutional and personal support that made this possible.

I would also like to acknowledge a small but memorable moment that shaped the first chapter of this thesis. At our initial supervisory meeting, when asked which industry I was interested in studying, I answered simply: robots. The response was immediate interesting, but far too broad. After a brief pause, and a glance around the room, the discussion narrowed quite literally to robotic vacuum cleaners. That casual moment set the foundation for Chapter 1. Not long afterwards, Amazon announced its proposed acquisition of iRobot, transforming an exploratory idea into a timely research question. As this sequence of events illustrates, chance often plays a role in research, but it tends to reward a prepared mind.

My heartfelt thanks go to my partner, Yanrong Zhang, whose love, understanding, and steadfast support sustained me through every stage of this journey. We met in 2022, when she was pursuing her own PhD in Chemistry at the University of Groningen, and our relationship developed in parallel with two rather different doctoral projects—one concerned with markets and mergers, the other with molecules and reactions. Over the years, this meant frequent travel between Norwich and Groningen, shared periods of intense work, and an ongoing exchange of ideas across disciplinary boundaries. I am grateful to Wesley R Browne at the University of Groningen for kindly providing me with a desk during my visits, which allowed me to work alongside her research group. Along the way, Yanrong acquired a working familiarity with mergers and acquisitions, while I gained a basic appreciation of inorganic chemistry—suggesting that even an economist and a chemist can find common ground. More importantly, we encouraged one another through the inevitable uncertainties of PhD life, and her companionship, patience, and perspective have been an enduring source of strength throughout this process. I am equally grateful to my parents for their enduring encouragement and for instilling in me the values that made this work possible. Although they remained in China throughout my doctoral studies, the distance never diminished their support: our regular weekly phone calls, constant care, and unwavering confidence in me provided reassurance and motivation at moments when perseverance mattered most. Their quiet sacrifices and steadfast belief have been a continual source of strength during my PhD journey.

I would also like to thank the School of Economics at the University of East Anglia, as well as colleagues across the Centre for Competition Policy, for providing a stimulating and interdisciplinary research environment. The Centre offered valuable opportunities to engage with researchers from different fields and to participate in interdisciplinary projects, which significantly broadened my perspective on competition policy and empirical research. I am particularly grateful to Farasat Bokhari, Data Editor at *Economic Inquiry*, for offering me the opportunity to work as a research assistant data editing. This experience substantially strengthened my understanding of research transparency and reproducibility, and deepened my practical knowledge of managing complex data archives and replication materials for conditionally accepted academic papers. I am also sincerely appreciative of the administrative and academic staff whose professionalism and support made my time at UEA both productive and rewarding.

Finally, this thesis made limited use of the generative artificial intelligence tool ChatGPT (OpenAI, GPT-4 and GPT-5 series, accessed during 2024–2025). The tool was used for language-related assistance (including grammar, clarity, and academic style) and for programming support, such as code formatting, syntax checking, debugging suggestions, and efficiency or readability improvements in scripts written by the author. All substantive intellectual content, methodological choices, analyses, and conclusions are entirely the responsibility of the author.

Chapter 1

Introduction

Digital markets are environments in which search, matching, and transactions are intermediated by large online platforms rather than by traditional retail or distribution channels. In these settings, competitive conditions are shaped not only by prices, but also by the collection and use of data, the ranking and visibility of products, personalised recommendations, the design of digital interfaces, alongside the fees charged by the platforms. These features give platforms substantial influence over how consumers perceive and access quality.

M&A in digital markets can modify competitive conditions through mechanisms that standard price-based tools only partially capture. Recent policy reports and regulatory initiatives—including the Stigler Center's report on digital platforms, the Organisation for Economic Co-operation and Development (OECD) Competition Assessment Toolkit, the United Kingdom (UK) Digital Markets Taskforce, the EU's Digital Markets Act, and Germany's Section 19a Gesetz gegen Wettbewerbsbeschränkungen (GWB)—stress that platforms may distort competition through algorithmic governance of information, visibility, and access, even in the absence of overt price increases. These conceptual challenges are reinforced by practical capacity constraints within competition authorities. Recent overviews of competition policy highlight that digital markets expose the limitations of the traditional antitrust toolkit and increase the need for more data-driven and forward-looking analytical approaches, particularly in market definition, evidentiality standards, and the treatment of non-price effects (Cowan et al. (2025)).

Digital merger cases require significant technical analysis – large volumes of data, sophisticated econometric evidence, and often complex algorithmic assessments – yet agencies operate under tight resource limits. The UK Competition and Markets Authority (CMA), for example, reported net expenditure of £122.6 million and an average staff of 1,127 Full-time equivalent (FTE)s in 2023/24.[1] The European Commission's Directorate-General for Competition (DG COMP) had €176.4 million in appropriations for 2024 and around 1,012 staff,[2] while the US Federal Trade Commission and Department of Justice Antitrust Division requested $430 million and $233.3 million respectively for Fiscal Year (FY) 2024.[3] These figures highlight a structural tension: the growing complexity of digital markets is not matched by a proportionate expansion in regulatory capacity, increasing the value of analytical tools that allow earlier, more targeted, and more efficient merger assessment.

Meanwhile, competition authorities increasingly recognise that non-price dimensions – innovation, product variety, quality, privacy, and sustainability – are central to welfare in digital markets. Yet existing empirical frameworks remain largely oriented toward environments with clear prices, stable product boundaries, and limited multi-sided interactions. Digital markets, by contrast, often feature zero monetary prices, rapid product turnover, and platform-mediated quality signals, making traditional Small but Significant Non-transitory Increase in Price (SSNIP)-based market definition ill-posed.

This thesis addresses these challenges through three self-contained but closely related essays. Together, they develop tools to conceptualise, measure, and empirically analyse non-price competitive effects of mergers in digital environments, with particular attention to platform-driven quality signals and to market definition in zero-price settings. Methodologically the chapters combine structural modelling, reduced-form event-study analysis, and modern NLP. Empirically, the thesis draws on two major digital markets: Amazon's marketplace for robotic vacuum cleaners and the mobile- app ecosystem on Google Play.

Across the three essays, the thesis is organised around three core research problems.

First, how can we model and quantify the competitive impact of *platform-driven*

[1] Competition and Markets Authority (2024), *Annual Report and Accounts 2023/24*, HC 7483, pp. 151, 210.

[2] European Commission (2023), *Draft General Budget of the EU for 2024*, Part II, Title 33, p. 422; European Commission (2023), *DG COMP Management Plan 2023*, p. 5.

[3] Federal Trade Commission (2023), *FY 2024 Congressional Budget Justification*, p. 1; US Department of Justice Antitrust Division (2023), *FY 2024 Congressional Budget Submission*, p. 5.

quality contributions in the context of non-horizontal mergers? When a dominant platform acquires a supplier, it may be able to tilt algorithmic ratings or rankings in favour of the acquired brand. This raises the question of how such subtle shifts in perceived quality translate into changes in prices, real quality choices, market shares, and welfare, and how merger policy should account for such non-price levers. The Amazon–iRobot non-consummated acquisition provides a salient motivating example: regulators explicitly worried that Amazon could favour iRobot through its control over search, recommendations, and review visibility. This thesis retrospectively assesses the complex impact of this merger.

Second, how do mergers and acquisitions in digital app markets affect *non-price firm behaviour* such as entry, product variety, innovation (app updates), and monetisation strategies? Using a large panel of Google Play apps and developers, the thesis asks whether M&A tend to expand or contract app portfolios, whether they reallocate effort from entry to maintenance, and how patterns differ between horizontal and non-horizontal deals, and between mergers and acquisitions.

Third, in digital environments where user-side prices are often zero, how can we construct *economically meaningful market definitions* that allow standard merger screens (e.g. HHI thresholds, change in concentration (Change in the Herfindahl–Hirschman Index (ΔHHI))) to be applied, and what do such screens reveal about the non-price consequences of consolidation? The thesis studies whether markets formed from textual product descriptions and change logs 'text-defined markets' can serve as a basis for concentration metrics, and how M&A within these markets reshape exit, update activity, and prices.

Taken together, these questions are intended to illuminate how digital mergers operate through quality, innovation, and algorithmic visibility rather than solely through prices, and to develop tools that make such channels empirically tractable for policy.

The thesis speaks to three main strands of literature: (i) merger analysis with prices and quality, (ii) non-price effects of mergers in digital markets, and (iii) market definition and measurement in data-rich but price-poor environments. Each strand has developed rapidly, but together they leave gaps that motivate the three chapters of the thesis.

A large literature studies the effects of mergers on prices and, increasingly, on product quality, using both reduced-form and structural approaches. Seminal structural work by Berry et al. (1995); Nevo (2001) develops RCL models that allow for heterogeneous consumer preferences and joint estimation of demand and supply, providing a standard platform for merger simulation. Subsequent contributions extend this framework by incorporating endogenous product quality and richer supply-side behaviour. For example, Crawford et al. (2019) embed quality choices directly into firms' cost functions, so that quality becomes a strategic decision variable that interacts with pricing, while work on multi-product oligopoly highlights how ownership consolidation affects portfolio composition and pricing across related products (e.g. Nocke and Schutz, 2018). Building on these foundations, recent merger theory explicitly links ownership changes to firms' incentives to invest in quality and innovation: Federico et al. (2018) and Motta and Tarantino (2021) show that when firms can adjust quality, mergers tend to weaken quality incentives unless sufficiently strong efficiency or innovation synergies are present, so that improvements in consumer surplus are not guaranteed. Taken together, this literature clarifies the conditions under which mergers may distort quality provision relative to the pre-merger benchmark.

By contrast, much less is known about *non-horizontal* mergers in which a digital platform acquires a supplier and can influence competition through algorithmic quality signals, recommendation systems, or ranking bias. A growing empirical literature on online reviews and rating manipulation documents systematic distortions in rating systems and their consequences for consumer information and market outcomes (Mayzlin et al., 2014; Luca and Zervas, 2016). Related work on search bias and pay-for-placement shows that integrated or advertising-funded platforms have incentives to tilt rankings in favour of affiliated or paying products (Hagiu and Jullien, 2011; de Cornière and Taylor, 2014). However, these mechanisms are typically analysed outside of a structural merger framework: very few models jointly specify demand, pricing, and platform-driven quality contributions in order to evaluate the competitive effects of a non-horizontal acquisition.

The gap, which **Chapter 2** addresses, lies in integrating platform-driven quality contributions into a structural merger model. The chapter develops a simple duopoly framework in which a platform can allocate 'quality contributions' (for example via ratings or visibility) across sellers, and then estimates a joint demand–supply system for the US robotic vacuum market under alternative ownership structures. This allows the analysis of how small shifts in perceived ratings under vertical integration translate into changes in prices, real quality choices, market shares, and welfare.

A second strand of work studies mergers in digital and high-tech markets with an emphasis on non-price outcomes such as innovation, product variety, privacy, and market dynamism. Classic contributions by Schumpeter J A (1942) and Arrow (1972) highlight the tension between market power, dynamic efficiency, and incentives to invest in new products and technologies. These ideas continue to inform recent analyses of mergers in environments where intangible assets and platform design shape productivity and competitive pressure. Empirically, Haucap et al. (2019) combine a formal model with evidence from pharmaceutical mergers to study how consolidation affects innovation, showing that changes in market structure can substantially alter both the level and the allocation of Research and Development (R&D) effort across overlapping technology fields. More recently, Haucap and Stiebale (2023) survey and extend the evidence on *non-price* effects of mergers and acquisitions, documenting how consolidation can influence innovation, product variety, and other quality dimensions even when prices remain relatively stable. Complementing this evidence, Bennato et al. (2021) show that mergers in the hard-disk-drive industry generate heterogeneous innovation responses, with different indicators—patents, R&D activity, and technological performance—often pointing in different directions. This reinforces the difficulty of assessing non-price merger effects using any single innovation metric. In the context of digital industries, Cabral (2025) calibrates a model of startup innovation, big-tech acquisitions, and merger review, highlighting how tighter merger control can reshape dynamic innovation incentives even when short-run price effects are muted. This literature systematically broadens the focus of merger control beyond short-run price effects, but it is still largely based on traditional industries and on outcomes measured via patents and R&D expenditures. By contrast, digital platform settings raise additional non-price concerns related to data, ranking, and the algorithmic governance of quality, for which the available evidence remains comparatively limited.

Within digital platform markets, researchers have begun to document how large platforms' acquisitions affect innovation, market concentration, and privacy. Policy reports by the European Commission and related initiatives express concerns that a permissive merger regime may fail to capture the elimination of potential competition or the role of data as a source of market power (European Commission, 2019; Stigler Committee on Digital Platforms, 2019). Affeldt and Kesler (2021) provide evidence that some Big Tech acquisitions are associated with weaker innovation efforts among competing apps, while Kesler et al. (2019) find that higher concentration on app platforms is correlated with more intensive data collection, raising privacy concerns. These contributions highlight important non-price channels but typically focus on a small number of very large acquirers (GAFAM) and on outcomes such as patents, entry, or privacy permissions.

Two gaps emerge from this literature. First, we have relatively little systematic evidence on firm-level portfolio decisions in digital markets within a unified empirical framework. Second, existing work often does not differentiate sharply between mergers and acquisitions, or between horizontal and non-horizontal deals, within a single digital ecosystem. **Chapter 3** speaks directly to these gaps by providing a detailed empirical study of non-price developer behaviour around 1,387 M&A in the Google Play app economy. Using panel data on apps and developers, the chapter distinguishes mergers from acquisitions and horizontal from non-horizontal events, and traces how new launches, external app purchases, removals, sales, updates, and monetisation strategies evolve before and after M&A at both the developer and app level. This evidence complements existing work on large platforms by documenting systematic patterns of portfolio adjustment and innovation for a broader set of app developers within a single digital ecosystem.

A third body of work concerns market definition and merger screening. In traditional settings, demand estimates and diversion ratios underpin the delineation of relevant markets and the calculation of concentration measures such as HHI, which in turn feed into merger screens and simulation exercises. Platform theory emphasises cross-side effects and net price/quality trade-offs in multi-sided markets (Rochet and Tirole, 2003; Armstrong, 2006). However, when user-side prices are zero, SSNIP tests become ill-posed on a single side, and quantity or ad-price changes may not provide a sufficient basis for defining markets relevant for merger assessment.

At the same time, textual descriptions of products and change logs are now abundant. In other domains, text-based similarity measures have been used to study product or industry spaces, for example in empirical finance and related applications

(Hoberg and Phillips, 2016; Gentzkow and Shapiro, 2019). There are also emerging applications of NLP tools to competition-policy questions, including the use of embeddings or topic models to characterise product overlap or proximity. Despite this, there are only limited examples where such text-based tools are translated into operational market definitions that can be used together with standard concentration thresholds and modern DiD estimators to study the consequences of mergers.

In particular, existing work provides limited evidence on (i) whether markets formed from product text can support policy-relevant concentration screens and (ii) which non-price outcomes such screens help predict in zero-price digital settings. Much of the current debate on digital merger control still relies on conventional, category-based market definitions that may not fully reflect how consumers perceive similarity in rich product spaces.

Chapter 4 addresses this gap by constructing 'text-defined markets' from app descriptions and update notes on Google Play and then combining these markets with EU- and US-style concentration screens and staggered DiD designs. The chapter uses these text-defined markets to study how mergers within a given market cluster affect exit, update activity, and prices, and compares deals that would and would not trigger standard concentration thresholds. In doing so, it provides an applied example of how NLP-based market definitions can be integrated with non-price outcome measures in a setting where user prices are frequently zero.

Across these three strands, the thesis argues that understanding mergers in digital markets requires moving beyond purely price-based paradigms. Platform-driven quality contributions, app-level portfolio decisions, and text-based market boundaries all play a role in shaping competitive outcomes, and the three chapters contribute to these literatures by providing theory, structural evidence, and empirical tools that make these channels more transparent for merger assessment.

The thesis is organised into three core chapters, each written as a stand-alone paper, followed by a general thesis conclusion.

Chapter 2 – 'Cleaning up? The effect of quality contribution in Amazon's attempted iRobot acquisition' develops the theoretical framework of platform-driven quality contributions and implements the structural RCL model for the US robotic vacuum market, using it to conduct counterfactual simulations of rating privileges under vertical integration.

Chapter 3 – 'Non-price effects of M&A in the app market' provides a reduced-form analysis of mergers and acquisitions in the Google Play ecosystem, using matched timing-of-events methods to study how M&A affect app portfolios, update activity, monetisation, and strategies at both the developer and app level.

Chapter 4 – 'Text-defined Markets: An NLP framework for antitrust with evidence from the app economy' proposes and implements the text-defined markets methodology, combines it with EU- and US-style merger screens, and uses staggered DiD estimators to examine the effects of M&A on exit, update activity, and prices within text-defined markets.

Overall, the thesis argues that understanding mergers in digital markets requires moving beyond traditional price-centric paradigms. Platform control over quality signals, app-level portfolio decisions, and market boundaries defined through text all play crucial roles in shaping competitive outcomes. By combining structural modelling, large-scale empirical analysis, and modern NLP-based tools, the thesis offers both conceptual and practical contributions to the design of merger policy in data-rich, algorithmically mediated markets.

Chapter 2

Cleaning up? The effect of quality contribution in Amazon's attempted iRobot acquisition[1]

Amazon's proposed acquisition of iRobot is analysed using an intuitive model of platform-driven quality allocation, demonstrating that platform ownership can amplify quality asymmetries—especially after acquiring a high-quality seller. Structural estimates based on US robotic vacuum data reveal that even modest, algorithmically driven boosts to iRobot's ratings post-merger can substantially increase both firm and platform profits, as well as shift market shares. These findings underscore the competitive importance of algorithmic quality contributions and indicate that merger policy should also account these non-price effects in digital markets.

[1]This chapter is based on joint work with Franco Mariuzzo. The initial theoretical framework was developed by Franco Mariuzzo. I implemented the model computationally, extended the original framework, and developed the full empirical analysis based on this theoretical structure. I collected and constructed all datasets and wrote the first complete draft of the chapter. Franco Mariuzzo provided guidance and comments throughout the project, including feedback on earlier drafts. All remaining errors are our own. We thank participants at the NIE Early Career Researchers' Symposium (Nottingham University Business School), the *Advances in Industrial Organization and Competition Policy* workshop (Reading University), the 3rd Economics PhD Conference 2025 (Essex Economics School, Colchester), the internal seminar at the School of Economics (UEA), and the 19th International Conference on Competition and Regulation (MINOA Palace Resort, Chania, Crete) for helpful comments and discussions. This manuscript benefited from the use of AI-assisted writing tools for grammar and style improvements, with all intellectual content and interpretations remaining the sole responsibility of the authors.

2.1 Introduction

The rise of digital platforms has fundamentally transformed competition, introducing new challenges for traditional merger control and antitrust frameworks. Unlike conventional horizontal mergers, platform acquisitions enable dominant firms to exert influence through non-price mechanisms, such as algorithmically-driven quality signals. This shift has heightened regulatory scrutiny, particularly regarding platforms' capacity to shape competitive outcomes through subtle interventions in product ratings and rankings, among other factors.

Amazon's proposed acquisition of iRobot—a leading manufacturer of robotic vacuum cleaners—serves as an example of these emerging challenges. Regulators, notably the European Commission,[2] and the US Federal Trade Commission (FTC),[3] have expressed concerns that Amazon's market role might enable it to strategically adjust product ratings, potentially disadvantaging rival products and entrenching iRobot's market position. Such platform-driven quality contributions could significantly alter competitive interactions, consumer perceptions, and market outcomes. Consequently, merger policy is increasingly adapting to account for these non-price effects, recognizing their potential to create important quality-driven competitive distortions.

This episode is emblematic of broader, ongoing policy debates. Seminal frameworks such as the US Stigler Center's Final Report on Digital Platforms (2019),[4] the OECD Competition Assessment Toolkit (2019),[5] and the UK Digital Markets Taskforce Report (2020),[6] all document how dominant platforms can distort competition by leveraging data, rankings, and infrastructure. Legislative and regulatory responses have since accelerated, from the US Department of Justice (DOJ)/FTC Vertical Merger Guidelines (2020),[7] and the EU's Digital Markets Act (2022),[8] to Germany's Section 19a of the Act Against Restraints of Competition, which pioneers algorithmic dominance criteria.[9] Empirical evidence, such as that

[2] `https://ec.europa.eu/commission/presscorner/detail/en/STATEMENT_24_521`
[3] `https://tinyurl.com/ms5ycfny`
[4] `https://tinyurl.com/mrxurx4p`
[5] `https://www.oecd.org/daf/competition/assessment-toolkit.htm`
[6] `https://www.gov.uk/cma-cases/digital-markets-taskforce`
[7] `https://tinyurl.com/5a3u835u`
[8] `https://tinyurl.com/mrxwr8an`
[9] `https://tinyurl.com/5eue6jat`

provided by Australia's Digital Platform Services Inquiry (2020-2025),[10] highlights how self-preferencing by platforms leads to real consumer harm and competitive distortions. Together, these developments underscore that market power in digital markets is increasingly exercised through non-price mechanisms—especially control over information, rankings, and ecosystem design.

This paper develops an intuitive theoretical setting and a structural modelling counterpart to retrospectively assess the Amazon-iRobot non-consummated merger. We present a seller duopoly model with platform-driven quality contribution, showing that a profit maximizing platform treats sellers equally when they are symmetric, but favors the advantaged seller under asymmetry. After acquisition, the platform systematically allocates more quality contribution to its owned target, regardless of the initial advantage.

Building on this, we empirically estimate a structural model capturing firms' strategic pricing and quality decisions by integrating both consumer demand and supply-side equations. Using detailed product-level data from Keepa.com and sales estimates from AMZScout,[11] we focus on the US robotic vacuum market. Our RCL demand framework addresses the endogenicity of price and quality-proxied by reviews rating-building on established structural demand literature (Berry et al., 1995; Nevo, 2001; Crawford et al., 2019). While data limitations prevent direct estimation of the platform's endogenous quality choices, we evaluate platform behaviour via counterfactual scenarios imposing exogenous quality contributions, thereby quantifying the impact of platform-driven quality bias in the hypothetical Amazon-iRobot context.

Our counterfactual analysis shows that even modest, platform-driven increases in perceived product quality can generate economically large equilibrium effects. Rating boosts of 0.10–0.50 stars—well within observed variation in online ratings-lead to several-hundred-percent increases in iRobot's market share and substantial gains in its stand-alone profits. In equilibrium, iRobot cuts prices sharply and increases endogenous quality, while rival firms raise prices modestly and also increase quality as market participation expands. Consumer surplus rises and the share-weighted average price across inside goods falls, highlighting perceived-quality contribution as

[10] `https://tinyurl.com/yr5knjvp`
[11] `https://amzscout.net/blog/amazon-reviews-importance-and-impact/`

a channel through which platforms can reshape prices, quality provision, and surplus allocation.

Our research contributes to the economic literature in several respects. First, it outlines how platforms can use quality contributions to influence competitive outcomes, thus extending conventional price-focused merger analyses. Second, we provide empirical evidence of the potency of product ratings as strategic competitive instruments (Chevalier and Mayzlin, 2006; Luca, 2016; Hollenbeck, 2018; Sun, 2011; Long and Liu, 2024). Finally, our findings provide policy-relevant evidence that underscores the necessity for merger reviews in digital markets to explicitly incorporate algorithmic and non-price competitive factors, aligning with the emerging regulatory changes that address these issues.

The remainder of the paper is organized as follows. Section 2.2 reviews the relevant literature on mergers, non-price effects, and platform quality contribution, and provides background on the robotic vacuum cleaner industry. Section 2.3 develops the theoretical model of platform-driven quality contribution and outlines the structural empirical framework. Section 2.4 describes the data, empirical implementation, and presents the main results, including counterfactual analyses of the proposed acquisition. Section 2.5 concludes with policy implications and directions for future research and refinements.

2.2 Background and literature review

2.2.1 Institutional and market background

2.2.1.1 The robotic vacuum cleaner industry

Robotic floor care has been a rapidly expanding consumer-durables category since the early 2000s. After Electrolux's prototype in the late 1990s and commercial launch in 2001, iRobot's *Roomba* (2002) delivered the first mass-market breakthrough at a \$200 price point.[12] Subsequent diffusion was driven by step-changes in navigation and mapping (e.g., vSLAM/LiDAR), on-board compute, and integration with the smart-home stack (e.g., Amazon Alexa voice control from

[12] `https://tinyurl.com/2nw2ra6f`.

2017), which raised perceived quality and reduced usage frictions.[13] Industry trackers document rapid scale-up-from roughly $12bn in 2021 toward $50.65bn by 2028-as premium features (auto-empty/auto-wash docks) migrated downmarket.[14] iRobot's early leadership—especially in the US-reflects first-mover advantages and proprietary navigation/cleaning know-how, later reinforced by capability and channel acquisitions: Evolution Robotics (2012, navigation/SLAM), Robopolis (2017, EU distribution), Root Robotics (2019, education/brand), and Aeris (2021, air purification).[15]

Table 2.1: iRobot M&A summary

Acquiree	**Acquiror**	**Announced**	**Acquiree description**
Evolution Robotics	iRobot	2012-09-17	Designing and constructing robots.
Sales On Demand Corp	iRobot	2016-11-21	Distribution business.
Robopolis	iRobot	2017-07-25	Supplier of robotics for household and educational uses.
Root Robotics	iRobot	2019-06-20	Learning system to teach coding from early childhood through adulthood.
Aeris Health	iRobot	2021-11-18	Smart indoor air-quality solutions.
iRobot	Amazon	2022-08-26	Designing and building behaviour-based AI robots.*

* Amazon announced its intention to acquire iRobot for $1.7 billion on August 26, 2022, but terminated the deal in January 2024 following opposition from the European Commission.
Notes: Source is Crunchbase. Table lists publicly announced mergers and acquisitions involving iRobot.

In 2022, Amazon announced a merger agreement with iRobot worth $1.7 billion, which is the fourth-largest in Amazon's M&A deals. It will strengthen Amazon's dominant position and create data security and privacy issues. There is a potential need to analyse M&A in this industry. We attempt to create a structural model to assess the prospective horizontal merger effect not just in the robotic vacuum industry, but also in other micro robotic industries.

2.2.1.2 Evolution of Amazon's product rating and review policies

Amazon's product review system has evolved significantly, reflecting the platform's increasing recognition of reviews as critical strategic instruments capable of influencing consumer perceptions and competitive market dynamics. Historically,

[13] https://tinyurl.com/2s48f95r.
[14] https://tinyurl.com/6tmjc5f2.
[15] https://tinyurl.com/3xwnzsa7; https://tinyurl.com/d3tv94p7; https://tinyurl.com/na9t33v6; https://tinyurl.com/4pkc99w4.

Amazon's policies have aimed both at enhancing the credibility of product ratings and at combating review contribution. A summary of key milestones in Amazon's review policy evolution is provided in Table 2.2.

Table 2.2: Evolution of amazon's product review and rating policies

Year	Policy Change	Description / Rationale
2007	Amazon Vine (invitation-only review programme)	Launch of the Vine programme, where selected reviewers receive free products to write reviews, aiming to promote credible and trustworthy feedback.
2015	Weighted Star Ratings (machine-learning model)	Amazon updates its star rating system using a machine-learning algorithm that assigns higher weights to recent, verified, and helpful reviews.[a]
2016	Ban on Incentivized Reviews	Amazon prohibits reviews written in exchange for incentives (e.g., free or discounted products), except through the Vine programme, to reduce biased or manipulative content.
2019	One-Tap Ratings Introduced	Introduction of a simplified rating feature allowing users to give star ratings without textual reviews-enhancing participation but reducing qualitative feedback.

[a] The weighting model prioritizes reviews that are more recent, verified, and marked as helpful by other users.
Sources: Wikipedia (2007); GeekWire (2015); TechCrunch (2016, 2019).

Amazon's policy evolution from a relatively transparent rating system to a highly curated, algorithmically weighted approach underscores its recognition of ratings as powerful strategic tools. The significant policy milestones—including the introduction of weighted ratings (2015) and the prohibition of incentivized reviews (2016)-demonstrate Amazon's acknowledgment of rating authenticity concerns and its commitment to maintaining the credibility of its review ecosystem. Yet, these policies simultaneously illustrate Amazon's increasing influence over perceived product quality through sophisticated, algorithmic mechanisms that can significantly shape consumer choice.

While there is no direct evidence that Amazon intentionally manipulates product ratings to disadvantage competitors, the documented history highlights Amazon's growing capacity to influence ratings indirectly through platform design choices. These policy shifts provide empirical grounding for our structural simulations, illustrating how subtle changes in the perceived quality signals managed by platforms could impact competitive outcomes significantly. Such indirect control over ratings becomes particularly relevant in analysing competitive implications in cases like Amazon's hypothetical acquisition of iRobot.

2.2.1.3 Seller behaviour, platform Influence, and vertical integration on Amazon

On Amazon, sellers influence product ratings and consumer perceptions not only through conventional attributes such as price and intrinsic product quality but also via a variety of strategic actions aimed at shaping consumer opinions. Legitimate practices commonly include enhancing customer service, offering targeted promotions, improving product packaging and branding, and soliciting genuine customer feedback through follow-up communications. However, alongside these acceptable methods, manipulative tactics have emerged and proliferated. Examples include incentivize positive reviews, employing fake review services, or engaging in review brushing-the practice of placing fake orders to artificially inflate review counts and ratings. Such activities significantly distort Amazon's reputation system, mislead consumers, and erode trust in the platform Gandhi et al. (2025).

To mitigate these manipulative tactics, Amazon continuously refines its moderation policies through both algorithmic and manual interventions. Specific mechanisms include automated detection and removal of fake or biased reviews, prioritization of verified purchase reviews, and elevation of reviews deemed particularly helpful by users. Despite these proactive measures, Amazon's enforcement remains imperfect. Manipulative strategies constantly evolve in response to platform rules, posing ongoing challenges to review authenticity and algorithmic detection methods.

Amazon's position within its own marketplace raises significant concerns about neutrality, as the firm operates simultaneously as a platform and a seller of competing products. Competition authorities, including the FTC and the European Commission's DG COMP, have scrutinized Amazon's practices, investigating allegations of preferential treatment for its own products. These concerns focus on potential advantages such as enhanced search result placements, exclusive access to valuable competitor sales data, and biased review contributions. Such practices could unfairly disadvantage independent sellers by limiting their market visibility and consumer appeal (European Commission (2020), Federal Trade Commission (2023a), and Federal Trade Commission (2023b)).

2.2.2 Literature review

2.2.2.1 Mergers and non-price effects

Due to the fast-changing nature of markets, mergers can significantly influence non-price dimensions, such as labour employment, sustainability, innovation, and product quality (Haucap and Stiebale (2023)). These effects are particularly critical in differentiated product markets, where consumer welfare depends on both price and quality improvements.

Mergers can have profound effects on labour market outcomes. Recent research has emphasized the role of labour market monopsony power, where increased concentration allows firms to suppress wages and reduce employee bargaining power. Marinescu and Hovenkamp (2019) highlight how antitrust policy, including merger scrutiny, should account for worker interests, not just consumer welfare. Empirical evidence from Prager and Schmitt (2021) demonstrates that higher regional concentration in hospital markets has led to lower wage growth for skilled workers, underscoring the labour implications of mergers in specific sectors. Regarding sustainability, the literature remains nascent but offers important insights. Recent studies, such as Liang et al. (2022), observe that mergers in polluting industries, such as energy, can sometimes foster green innovation when supported by subsidies and policy incentives. However, results remain mixed overall, with some evidence suggesting that mergers can reduce competition necessary to drive eco-innovation. Aghion et al. (2023) discuss the broader relationship between sustainability and competition, noting that higher competition levels often lead to improved environmental outcomes by pushing firms toward more efficient and sustainable practices.

Innovation remains one of the most studied non-price effects of mergers. Federico et al. (2018) and Motta and Tarantino (2021) find that while mergers can sometimes spur innovation by reducing duplication and fostering R&D synergies, they more often suppress innovation by diminishing competitive pressures. Haucap et al. (2019)

confirm these trends in the pharmaceutical industry, where mergers have been shown to reduce both the quantity and novelty of innovation efforts. Building on this literature, Valletti (2025) clarifies the 'Innovation Theory of Harm', emphasizing that absent synergies or spillovers, mergers are generally detrimental to consumer welfare because they both reduce innovation incentives and raise prices, reinforcing the need for merger control to scrutinize innovation and quality competition alongside prices. Product quality is another critical dimension. Haucap et al. (2019) note that mergers can lead to quality reductions due to decreased competitive pressure. However, certain vertical integrations, as discussed by Hansman et al. (2020), can lead to quality enhancements through improved supply chain coordination. A case study Bennato et al. (2021) demonstrates that the method of measuring innovation reveals varying effects —not always in the same direction—of mergers on innovation.

This study contributes to the literature by focusing on platform-driven quality contributions, where platforms like Amazon can influence quality, in our analysis measured by product ratings, and possibly distort competition in the market. By bridging demand-side modelling with endogenous quality adjustments, it provides insights into how platform dominance can impact quality, offering a new perspective on merger effects in digital markets. Additionally, the analysis of the robotic vacuum cleaner industry extends the literature by providing insights into non-price effects in emerging technology-driven markets, offering a nuanced understanding of the interaction between mergers, quality competition, and consumer benefits.

2.2.2.2 Ranking biases, platform contribution, and consumer welfare

Recent literature highlights significant concerns regarding how digital marketplaces strategically manipulate rankings and consumer feedback mechanisms (seen as proxies for product quality), influencing consumer decisions and competition. Unlike traditional markets, where price and intrinsic quality predominantly guide consumer choice, platform-driven markets introduce additional levers such as algorithmic ranking biases, review visibility adjustments, and sponsored advertising to shape consumer perceptions and choices.

A critical area of study is how algorithmic ranking biases affect competitive outcomes. For instance, de Cornière and Taylor (2014) illustrate that search engines and online marketplaces can intentionally introduce ranking biases, prioritizing specific sellers or products to enhance their own profitability, potentially undermining competition by limiting consumer exposure to alternatives. Choi and Mela (2019) extend this discussion by empirically demonstrating how platforms adjust rankings strategically to maximize advertising revenue from sellers, thus highlighting inherent conflicts of interest in platform operations. These manipulations directly impact consumer welfare by restricting access to potentially superior or more cost-effective alternatives.

Further research addresses the manipulation of online reviews and ratings, a phenomenon that significantly shapes consumer perceptions. Luca and Zervas (2016) document extensive evidence of review fraud, revealing systematic distortions in consumer perceptions caused by manipulated reviews, particularly prevalent in competitive sectors. Mayzlin et al. (2014) corroborate these findings, showing firms' strategic behaviours include inflating their own product ratings and deliberately lowering competitor ratings through deceptive reviews. Such practices undermine the reliability of online reputation systems and distort consumer decision-making processes. Gandhi et al. (2025) further dissect the implications of fake reviews, distinguishing between the misinformation and mistrust effects. Their findings show that misinformation, by inflating product ratings, leads consumers to make suboptimal choices, thereby decreasing overall consumer welfare. Simultaneously, mistrust in review systems, triggered by widespread manipulation, reduces the efficacy of reviews in signalling true product quality, further complicating consumer decision-making.

Platforms' influence also extends to sponsored advertising and recommendation algorithms. Armstrong and Zhou (2011) and Hagiu and Jullien (2011) find that platforms frequently prioritize products and services from sellers who pay higher fees, creating implicit pay-for-placement mechanisms that negatively influence organic competition. Long and Liu (2024) further document how platforms tolerate fake reviews and leverage sponsored advertising to skew market outcomes in their favour, reinforcing the centrality of algorithmic governance in digital marketplace. Zhu and Liu (2018) empirically demonstrate that platforms, such as Amazon, strategically enter into competition with complementers by targeting successful product spaces, thereby undermining third-party sellers' incentives and ultimately discouraging innovation and competitive entry.

Our study builds on these insights by empirically examining platform-driven rating contribution within the context of Amazon's proposed acquisition of iRobot. Specifically, our analysis demonstrates that even modest contributions of product ratings can lead to dramatic reallocation of market share and profits. Our counterfactual simulations indicate that minor artificial enhancements to iRobot's ratings significantly amplify its market position and profitability, while simultaneously disadvantaging rivals whose ratings are subtly diminished. These findings provide empirical evidence of the powerful non-price strategies platforms can employ.

By explicitly modelling consumer response and competitive interaction under manipulated product ratings, our research contributes a nuanced understanding of how rating contributions function as strategic non-price barriers. This enhanced understanding emphasizes the necessity for regulatory frameworks to incorporate oversight of platform-controlled rating and ranking systems, thus ensuring competitive fairness and protecting consumer welfare in digital markets.

2.2.2.3 Endogenous product choice

The analysis of mergers has predominantly focused on markets with single-product firms or where product quality is treated as an exogenous variable. These studies typically assume quantity competition (Perry and Porter (1985); Levin (1990); McAfee and Williams (1992)) or price competition with horizontally differentiated products (Deneckere and Davidson (1985)). However, the increasing complexity of modern markets demands a more nuanced approach, particularly in contexts where product quality and multi-product strategies are central. While much of the theoretical literature addresses these issues within the framework of horizontal mergers, less attention has been paid to non-horizontal mergers, such as platform-driven acquisitions.

In horizontal merger studies, Federico et al. (2018) and Motta and Tarantino (2021)) explore scenarios where firms can endogenously determine product quality. They demonstrate that significant innovation synergies are necessary for mergers to enhance consumer surplus, underlining the interplay between quality differentiation and welfare outcomes. Similarly, Nocke and Schutz (2018) examine multi-product firms using a nested CES/logit framework, showing that many classic results from single-product settings—such as mergers reducing consumer surplus absent synergies-remain valid in these cases. These contributions highlight the importance of incorporating quality and multi-product competition into merger analysis, even though they primarily address horizontal mergers.

This study extends the existing literature by shifting the focus to non-horizontal mergers, specifically Amazon's proposed acquisition of iRobot. This case illustrates the unique dynamics of platform-dominated markets where the acquiring firm can influence market outcomes not only through direct competition but also by leveraging platform-based mechanisms, such as ratings contribution, to shape demand and competitor quality. Unlike traditional merger settings, this acquisition underscores the strategic role of non-price effects and platform-mediated quality control in determining market power.

By endogenising product choice through quality differentiation, this paper investigates how non-horizontal mergers can affect market structure, consumer welfare, and competition in markets. The findings reveal mechanisms distinct from those observed in horizontal mergers, offering new insights into the competitive dynamics of platform-based ecosystems and the broader implications for antitrust policy.

2.2.2.4 Empirical analysis of mergers on price and quality

Empirical studies have also sought to evaluate the impact of mergers on both price and quality using structural models. Ashenfelter et al. (2013) provide robust evidence of price increases following mergers in consumer goods markets, particularly when the merging firms hold significant pre-merger market shares. Focarelli and Panetta (2003) examine the banking industry and find that mergers can lead to long-term quality improvements despite short-term price increases. Chandra and Collard-Wexler (2009) analyse mergers in the newspaper industry-an archetypal two-sided market-and highlight mixed results, with some mergers improving content quality but others reducing diversity due to decreased competition for readers and advertisers. Miller and Weinberg (2017) examine the beer industry and find that horizontal mergers often lead to price increases. These findings illustrate the importance of analysing both dimensions simultaneously to capture the full welfare implications of mergers.

Structural approaches, such as the random-coefficients logit model developed by Berry et al. (1995) and extended by Nevo (2001), have become standard tools for estimating consumer demand in differentiated product markets. These models allow for heterogeneous consumer preferences and enable researchers to infer quality changes indirectly through shifts in demand and pricing dynamics. Crawford et al. (2019) further demonstrate how incorporating both demand and supply equations can reveal quality over-provision in equilibrium, emphasizing the importance of modelling firms' strategic quality choices alongside pricing.

2.3 Theoretical and empirical models

2.3.1 An illustrative theoretical duopoly model

Consider a market with two single-product firms ($i = 1, 2$) selling horizontally differentiated products via a digital platform. The platform can contribute to perceived quality signals, affecting demand for each product. We assume linear demand functions and strictly convex and independent cost functions.

Assumption 1 (Demand System)**.** *Demand for firm i's product is given by:*

$$q_i = \lambda_i + \beta(g_i + d_i) - \beta\rho(g_j + d_j) + \rho p_j - p_i, \qquad i \neq j. \tag{2.1}$$

Here, λ_i denotes firm i's stand-alone demand shifter; g_i is the component of product quality controlled by firm i; d_i is the component of quality determined by the platform; β represents consumer sensitivity to perceived quality; ρ indicates the degree of substitutability (assumed symmetric for both prices and qualities) between the two products; and p_i is the price. The subscript j denotes the competing firm (or product). The direct effect of a firm's own quality and platform quality control on its demand is governed by β; the cross-effect from the rival's quality is always strictly smaller, as it is scaled by ρ.

The two quality variables are central to the analysis and play distinct economic roles. g_i is the seller-chosen, intrinsic component of quality: the firm invests in genuine product improvement (durability, performance, design) and bears the quadratic cost $\frac{1}{2}g_i^2$. d_i is the platform-controlled component: it does not alter the physical product but shifts how favourably the platform presents or amplifies the product's quality signal, for example through ratings, review curation, or ranking and visibility. Consumers cannot separate the two and respond only to the combined perceived quality $g_i + d_i$ entering demand in Equation (2.1); in the empirical model both are captured through observed product ratings, which serve as a proxy for perceived quality. The distinction matters because the seller chooses g_i to maximize its own profit, whereas the platform chooses d_i in pursuit of platform objectives, so vertical integration changes who controls each lever and for whose benefit it is set.

Assumption 2 (Costs and platform fees)**.** *Each firm faces constant marginal cost of production c and quadratic quality cost $\frac{1}{2}g_i^2$. The platform charges a revenue share ϕ to the sellers' use of the platform.*

Assumption 3 (Parameter and variable restrictions)**.** *Throughout, we assume:*

$$\begin{aligned} &\beta,\, c \geq 0, \qquad 0 \leq (\phi,\ \rho) < 1, \\ &p_1,\, p_2 > 0, \qquad g_1,\, g_2 \geq 0, \\ &\lambda_1,\, \lambda_2 \in \mathbb{R}, \qquad d_1,\, d_2 \in \mathbb{R}. \end{aligned} \tag{2.2}$$

Timing

1. The platform sets quality contribution levels (d_1, d_2).

2. Firms (sellers) observe (d_1, d_2) and simultaneously choose (p_i, g_i).

2.3.1.1 Benchmark case: Equilibrium without acquisition

Firm problem: Each firm i chooses price and quality,[16] to maximize the profit:

$$\pi_i = (1-\phi)p_i q_i(\cdot) - c q_i(\cdot) - \frac{1}{2}g_i^2. \tag{2.3}$$

[16]From a practical perspective, firms would choose their level of investment rather than quality directly. However, to simplify the model and the notation, we assume that they choose quality and incur a quadratic investment cost. The same logic applies to the platform.

Platform problem: The platform chooses (d_1, d_2) to maximize:

$$\Pi_a = \phi\Big(p_1 q_1(\cdot) + p_2 q_2(\cdot)\Big) - \frac{1}{2}(d_1^2 + d_2^2). \tag{2.4}$$

Equilibrium characterization: Equilibrium prices, qualities, platform contribution, and quantities satisfy the following system of equations (with second-order conditions satisfied under the assumed functional form):

$$\begin{aligned} &\text{Firm } i\text{'s FOCs:} && \frac{\partial \pi_i}{\partial p_i} = 0, \quad \frac{\partial \pi_i}{\partial g_i} = 0, \\ &\text{Platform's FOCs:} && \frac{\partial \Pi_a}{\partial d_1} = 0, \quad \frac{\partial \Pi_a}{\partial d_2} = 0. \end{aligned} \tag{2.5}$$

Let $(p_1^*, p_2^*, g_1^*, g_2^*, d_1^*, d_2^*)$ denote the equilibrium solutions.

Equilibrium construction and notation. The equilibrium is obtained by backward induction.

Stage 2 (firms): Given the platform's contributions $d \equiv (d_1, d_2)$, firms 1 and 2 play a simultaneous game in prices and qualities. Let $\pi_i(p_i, g_i, p_j, g_j; d)$ denote firm i's profit. Best responses $B_i^p(p_j, g; d)$ and $B_i^g(p, g_j; d)$ are defined by solving the first-order conditions $\partial\pi_i/\partial p_i = 0$ and $\partial\pi_i/\partial g_i = 0$, respectively.

Solving the two firms' first-order conditions jointly yields the four *reduced-form* mappings

$$\Psi(d) \equiv \big\{p_1^*(d), g_1^*(d), p_2^*(d), g_2^*(d)\big\}, \tag{2.6}$$

which depend only on d and the model's primitives.

Stage 1 (platform): Anticipating $\Psi(d)$, the platform chooses d_1 and d_2 to maximize

$$\Pi_a(d) = \phi[p_1^*(d)\, q_1^*(d) + p_2^*(d)\, q_2^*(d)] - \tfrac{1}{2}\left(d_1^2 + d_2^2\right), \tag{2.7}$$

where $q_i^*(d)$ are demands evaluated at $\Psi(d)$. The optimal contributions $d^* = (d_1^*, d_2^*)$ solve $\partial\Pi_a/\partial d_k = 0$ $(k = 1, 2)$, with second-order conditions verified by the Hessian of Π_a with respect to d. The subgame-perfect equilibrium is the fixed point $\big(\Psi(d^*), d^*\big)$.

Proposition 1 (Benchmark situation: sellers' independent of the platform)**.** *In the absence of platform acquisition, there exists a unique equilibrium in which the platform applies quality contribution to all sellers. We analyse three scenarios-(i) symmetric* *sellers ($\lambda_1 = \lambda_2$), (ii) seller one more prominent ($\lambda_1 > \lambda_2$), and (iii) seller two more advantaged ($\lambda_1 < \lambda_2$).*

1. *When sellers are symmetric, the platform contributes equally and positively to both sellers.*

2. *When one seller is more prominent, the platform contributes more to that seller, resulting in a higher price and quality, as well as greater market share and profit.*

Proof. Explicit solutions are provided in Appendix 2.A.1, and numerical results are summarized in Table 2.3. □

2.3.1.2 Equilibrium with acquisition

Revised platform profit: Suppose the platform acquires firm one and internalizes its profit. The new platform (inclusive of firm one) profit function is:

$$\Pi_a = (p_1 - c)q_1(\cdot) - \frac{1}{2}g_1^2 + \phi p_2 q_2(\cdot) - \frac{1}{2}(d_1^2 + d_2^2). \tag{2.8}$$

Firm 2's profit remains unchanged.

Equilibrium characterization Equilibrium prices, qualities, platform contribution, and quantities satisfy the following system of equations:

$$\begin{aligned} &\text{Firm 2's FOCs:} && \frac{\partial \pi_2}{\partial p_2} = 0, \quad \frac{\partial \pi_2}{\partial g_2} = 0, \\ &\text{Platform's FOCs:} && \frac{\partial \Pi_a}{\partial p_1} = 0, \quad \frac{\partial \Pi_a}{\partial g_1} = 0, \quad \frac{\partial \Pi_a}{\partial d_1} = 0, \quad \frac{\partial \Pi_a}{\partial d_2} = 0. \end{aligned} \tag{2.9}$$

The demand system and cost structure remain as previously defined. Let $(\tilde{p}_1^*, \tilde{p}_2^*, \tilde{g}_1^*, \tilde{g}_2^*, \tilde{d}_1^*, \tilde{d}_2^*)$ denote the new equilibrium solutions.

Equilibrium construction. With the acquisition, the platform internalizes firm 1's margin in its objective function (while still earning fee revenue on seller 2). The timing and decision rights are unchanged: in Stage 1 the platform chooses rating contributions $d = (d_1, d_2)$; in Stage 2 each seller $i \in \{1, 2\}$ chooses (p_i, g_i) to maximize π_i given d. Accordingly, firm 2's best responses satisfy $\partial\pi_2/\partial p_2 = 0$ and $\partial\pi_2/\partial g_2 = 0$

(and analogously for firm 1). The platform then selects d to maximize its modified objective $\tilde{\Pi}_a\big(d;\, p_1^*(d), p_2^*(d), g_1^*(d), g_2^*(d)\big)$. The resulting equilibrium is therefore

$$(\tilde{p}_1^*, \tilde{p}_2^*, \tilde{g}_1^*, \tilde{g}_2^*, \tilde{d}_1^*, \tilde{d}_2^*) = \big(p_1^*(\tilde{d}^*), p_2^*(\tilde{d}^*), g_1^*(\tilde{d}^*), g_2^*(\tilde{d}^*), \tilde{d}^*\big), \tag{2.10}$$

where tildes indicate that only the platform's objective (and thus $\tilde{d}^*$) differs from the benchmark.

Proof. Explicit solutions are provided in Appendix 2.A.1, and numerical results are summarized in Table 2.3. □

Table 2.3: Summary of parameter settings, equilibrium outcomes, and market results across scenarios

Scenario	**State**	β	c	λ_1	λ_2	ϕ	ρ	d_1	d_2	p_1 / p_2	g_1 / g_2	s_1 / s_2	π_1 / π_2	Π_a	CS
1 (Symmetric)	Pre	0.70	1.00	1.35	1.35	0.15	0.70	0.08	0.08	2.04 / 2.04	0.51 / 0.51	0.50 / 0.50	0.50 / 0.50	0.52	2.48
	Post	0.70	1.00	1.35	1.35	0.15	0.70	0.77	-0.32	2.39 / 1.74	0.72 / 0.33	0.68 / 0.32	— / 0.21	1.23	2.70
2 ($\lambda_1 > \lambda_2$)	Pre	0.70	1.00	1.50	1.20	0.15	0.70	0.09	0.06	2.12 / 1.95	0.56 / 0.46	0.55 / 0.45	0.60 / 0.41	0.52	2.83
	Post	0.70	1.00	1.50	1.20	0.15	0.70	0.86	-0.38	2.51 / 1.62	0.79 / 0.27	0.75 / 0.25	— / 0.14	1.37	2.76
3 ($\lambda_2 > \lambda_1$)	Pre	0.70	1.00	1.20	1.50	0.15	0.70	0.06	0.09	1.95 / 2.12	0.46 / 0.56	0.45 / 0.55	0.41 / 0.60	0.52	2.48
	Post	0.70	1.00	1.20	1.50	0.15	0.70	0.69	-0.27	2.27 / 1.85	0.65 / 0.40	0.62 / 0.38	— / 0.31	1.10	2.64

Notes: Each scenario column shows parameter values, platform contribution (d_i), equilibrium prices (p_i), equilibrium quality investments (g_i), market shares (s_i), seller profits (π_i), platform profit (Π_a), and consumer surplus (CS). β = demand sensitivity to quality; c = marginal cost; λ_i = demand shifter for firm i; ϕ = platform commission rate; ρ = cross-price/quality interaction; d_i = equilibrium quality contribution for firm i; p_i = equilibrium price; g_i = equilibrium quality investment; s_i = market share; π_i = profit; Π_a = platform profit; CS = aggregate consumer surplus. 'Pre' refers to the benchmark (no acquisition), 'Post' to the platform's acquisition of firm 1. For 'Post' rows, if a cell is marked '—', the value is not reported in the original table. Scenario 1: symmetric firms ($\lambda_1 = \lambda_2$); Scenario 2: firm 1 is stronger ($\lambda_1 > \lambda_2$); Scenario 3: firm 2 is stronger ($\lambda_2 > \lambda_1$).

Discussion: Table 2.3 summarizes equilibrium outcomes across three core scenarios that vary only in the degree of symmetry between sellers' baseline demand. In the symmetric baseline, where both sellers have equal intrinsic appeal ($\lambda_1 = \lambda_2$), pre-acquisition equilibrium is balanced: prices, quality, and market shares are all symmetric. Following acquisition, however, the platform's incentives realign to prioritize the owned seller. This results in a sharp increase in both price and quality for the acquired product, while the independent rival's market position is weakened-its price and quality both fall, and its market share shrinks. The post-acquisition reallocation is reflected in a substantial rise in platform profit, and, in this scenario, consumer surplus also increases as consumers benefit from higher average quality despite higher prices.

The second scenario introduces asymmetry, granting seller one a stronger baseline brand ($\lambda_1 > \lambda_2$). Prior to acquisition, seller one leads in price, quality, and market share. Post-acquisition, the platform's contribution and strategic behaviour further amplify this advantage. The acquired seller's price and quality both increase sharply, while the independent seller's outcomes are suppressed even further. This results in a dramatic shift in market shares and platform profit. Notably, however, this scenario demonstrates that acquisition does not always benefit consumers: in the case where pre-existing asymmetry is large and platform intervention is aggressive, consumer surplus may actually decrease after acquisition. Here, the negative effect of higher prices outweighs the gain from quality improvements, indicating a potential welfare loss and raising important policy considerations regarding the limits of platform integration.

In the third scenario, the initial asymmetry favours seller two ($\lambda_2 > \lambda_1$), and the pre-acquisition market is correspondingly tilted. Post-acquisition, the platform's behaviour again favours the owned seller, closing the initial gap and in some cases reversing it-seller one overtakes seller two in both price and quality. As in the symmetric case, both platform profit and consumer surplus rise after the acquisition, with the greatest gains accruing in scenarios where initial asymmetry is not overwhelming.

Taken together, these scenarios illustrate that the effect of platform acquisition is highly sensitive to underlying demand asymmetries. Across all cases, the platform leverages its ability to manipulate perceived quality and pricing, consistently shifting market share and profit toward the owned product. While consumer surplus typically rises-driven by improved average quality and intensified competition for visibility-there are clear instances where consumer welfare is reduced by post-acquisition price increases. These results highlight that platform acquisitions may generate both pro- and anti-competitive effects, depending on the market's initial structure, and underscore the need for nuanced antitrust scrutiny in digital marketplaces.

2.3.2 Econometric model

2.3.2.1 Pre-acquisition

In this chapter, it is useful to distinguish between two conceptually different components of product quality. First, g_{jt} denotes the seller-controlled component of quality, such as design, durability, suction performance, reliability, or customer support. This is the quality dimension the firm can improve through investment. Second, d_{jt} denotes the platform-controlled component of perceived quality, which operates through visibility and information design, for example by affecting how reviews are displayed, weighted, or surfaced to consumers. Consumers respond to the combination of these two components because both shape the product's perceived attractiveness. In the empirical implementation, observed ratings are treated as a proxy for perceived quality, capturing both seller-generated quality and any platform-mediated contribution to how that quality is presented.

Initially, the Amazon marketplace comprises multiple independent sellers (firms), each independently determining product-specific attributes such as the price of a product j, (p_j), and quality, (g_j). This quality decision directly influences product appeal and subsequent consumer reviews, encompassing both tangible attributes (e.g., durability, design) and intangible elements (e.g., brand reputation, customer support). Quality choices typically represent a strategic trade-off between cost and consumer appeal.

Additionally, the platform selects actions (d_j) affecting product visibility and review presentation, including algorithmic curation of reviews, adjustments in search rankings, and validation of review authenticity. The platform's actions are designed to foster platform integrity, maintain consumer trust, and maximize long-term engagement. The platform charges a commission rate ϕ on firms' revenues.

Each firm (seller) f chooses simultaneously the prices and qualities of the subset of products they control, to maximize every period t their profit

$$\max_{\{\boldsymbol{p}_{ft},\boldsymbol{g}_{ft}\}} \Pi_{ft} = \sum_{k\in\mathcal{J}_{ft}} \Big((1-\phi)p_{kt}\, s_{kt}(\mathbf{p}_t,\mathbf{g}_t;\mathbf{d}_t) \;-\; C_{kt}\big(g_{kt},\; s_{kt}(\mathbf{p}_t,\mathbf{g}_t;\mathbf{d}_t)\big)\Big) I_t, \quad (2.11)$$

where $s_{kt}(\cdot)$ is the market share of product k, and I_t denotes the market size in period t, that is, the total number of units sold in the market in that period rather than the number of consumers in the economy. Because demand is modelled in units, total unit sales of product k are therefore $q_{kt} = s_{kt}I_t$. It is worth emphasizing that $(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t)$ represent vectors encompassing all products marketed in period t. The cost function is modelled as $C_{kt}\big(g_{kt},\, s_{kt}(\mathbf{p}_t,\mathbf{g}_t;\mathbf{d}_t)\big)$, with (constant with respect to quantity) marginal cost of production and distribution:

$$\frac{\partial C_{kt}}{\partial s_{kt}} = c_{kt}(g_{kt}). \quad (2.12)$$

Subsequently, the platform selects its contribution to the quality effort across all firms by maximizing profit, subject to the per-user homogeneous cost function $K(\cdot)$.

$$\max_{\{\mathbf{d}_t\}} \Pi_{at} = \sum_{k\in\mathcal{J}_t} \big(\phi p_{kt}\, s_{kt}(\mathbf{d}_t) - K(d_{kt})\big) I_t. \quad (2.13)$$

The firms' profit maximization yields the following first-order conditions.

Price FOCs

$$\frac{\partial \Pi_{ft}}{\partial p_{jt}} = (1-\phi)s_{jt}(\cdot) + \sum_{k\in\mathcal{J}_t} \Big((1-\phi)p_{kt} - c_{kt}(\cdot)\Big) \frac{\partial s_{kt}(\cdot)}{\partial p_{jt}} \;=\; 0$$

$$\Longrightarrow \quad \sum_{k\in\mathcal{J}_t} \Big(1-\phi)p_{kt} - c_{kt}(\cdot)\Big) \Big(-\frac{\partial s_{kt}(\cdot)}{\partial p_{jt}}\Big) = (1-\phi)s_{jt}(\cdot). \quad (2.14)$$

In matrix form, let $\mathbf{s}_t, \mathbf{p}_t, \mathbf{c}_t$ be vectors and define the $\{0,1\}$ ownership matrix $\mathbf{H}_t(\cdot)$, to have the adjusted Hessian,

$$\boldsymbol{\Delta}_t^p(\cdot) \;=\; \mathbf{H}_t(\cdot) \circ \Big(-\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{p}_t}\Big). \quad (2.15)$$

Then,

$$\boldsymbol{\Delta}_t^p(\cdot)\Big((1-\phi)\mathbf{p}_t - \mathbf{c}_t(\cdot)\Big) \;=\; (1-\phi)\mathbf{s}_t(\cdot),$$

so that markups can be expressed as:

$$(1-\phi)\mathbf{p}_t - \mathbf{c}_t(\cdot) \;=\; (1-\phi)\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{s}_t(\cdot). \tag{2.16}$$

Quality FOCs

$$\frac{\partial \Pi_{ft}}{\partial g_{jt}} = \sum_{k\in\mathcal{J}_{ft}} \Big((1-\phi)p_{kt} - c_{kt}\Big)\frac{\partial s_{kt}(\cdot)}{\partial g_{jt}} - \frac{\partial C_{jt}(\cdot)}{\partial g_{jt}} \;=\; 0$$

$$\Longrightarrow \quad \sum_{k\in\mathcal{J}_{ft}} \Big((1-\phi)p_{kt} - c_{kt}\Big)\frac{\partial s_{kt}}{\partial g_{jt}} = \frac{\partial C_{jt}(\cdot)}{\partial g_{jt}}. \tag{2.17}$$

In parallel to above, define

$$\boldsymbol{\Delta}_t^g(\cdot) \;=\; \mathbf{H}_t(\cdot)\circ\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{g}_t}, \tag{2.18}$$

so

$$\boldsymbol{\Delta}_t^g(\cdot)\Big((1-\phi)\mathbf{p}_t - \mathbf{c}_t(\cdot)\Big) \;=\; \mathbf{c}_t^g(\cdot),$$

where $\mathbf{c}_t^g(\cdot) = \frac{\partial \mathbf{C_t}(\cdot)}{\partial \mathbf{g}_t}$ is the vector of marginal quality costs.

In this instance markups satisfy:

$$(1-\phi)\mathbf{p}_t - \mathbf{c}_t(\cdot) \;=\; \boldsymbol{\Delta}_t^g(\cdot)^{-1}\,\mathbf{c}_t^g(\cdot). \tag{2.19}$$

2.3.2.2 Post acquisition

When Amazon acquires an independent firm (f), the platform gains integrated control over product pricing, intrinsic product quality, and platform-mediated actions. This vertical integration enables Amazon to internalize the synergy between product quality and the strategic visibility of product reviews. For example, Amazon could coordinate higher quality investments with platform-driven review curation to enhance perceived product quality and prominence, reinforcing the value derived from these investments.

Simultaneously, Amazon retains control over platform-mediated actions for competing third-party sellers (d_{-f}). This dual position risks neutrality, as subtle alterations—such as the selective filtering of positive competitor reviews or the de-ranking of rival products-may systematically advantage Amazon's own products. Such behaviour may result from intentional strategic decisions or inadvertently emerge from inherent conflicts within Amazon's dual-role incentive structure.

Thus, post-acquisition, it is analytically coherent and realistic to model seller-controlled product quality as $\{g_{jt} : j \in \mathcal{J}_{ft}\}$, while platform-mediated actions remain product-specific and are represented by $\{d_{jt}\}$. This integrated framework captures the complex and consequential effects inherent in platforms operating simultaneously as neutral marketplaces and market participants, highlighting key competitive and regulatory implications.

Importantly, platform contribution remains product-specific after the acquisition: d_{jt} is chosen separately for each product j, not uniformly at the firm level. The acquisition changes the platform's incentives when choosing the vector $\mathbf{d}_t$, but it does not change the level at which d is defined.

There is a single platform (Amazon), and we use the subscript a purely to index platform-level quantities (for example, platform profit Π_{at}) and to distinguish them from seller-level quantities. After the acquisition, this platform owns firm f and therefore internalizes that firm's profits. The relevant profit maximization problems become:

$$\begin{aligned}\max_{\{p_{ft},\,g_{ft}\}} \Pi_{at} &= \sum_{k \in \mathcal{J}_{ft}} \Big(p_{kt}\, s_{kt}(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t) - C_{kt}\big(g_{kt},\, s_{kt}(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t)\big) - K(d_{kt})\Big) I_t + \\ &\quad + \sum_{l \notin \mathcal{J}_{ft}} \left(\phi p_{lt}\, s_{lt}(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t) - K(d_{lt})\right) I_t \\ \max_{\{p_{f't},\,g_{f't}\}} \Pi_{f't} &= \sum_{l \in \mathcal{J}_{f't}} \Big((1-\phi)\, p_{lt}\, s_{lt}(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t) - C_{lt}\big(g_{lt},\, s_{lt}(\mathbf{p}_t, \mathbf{g}_t; \mathbf{d}_t)\big)\Big) I_t, \quad f' \neq f.\end{aligned} \tag{2.20}$$

Here, d_{jt} denotes the platform's contribution to product j's perceived quality (e.g., boosting review visibility or ranking). This action incurs a homogeneous per-user convex cost $K(d_{jt})$, borne entirely by the platform. We assume $K(\cdot)$ is additive, and symmetric across products and increasing in magnitude, ensuring interior solutions.

Price FOCs

$$
\begin{aligned}
\text{(a) If } j \in \mathcal{J}_f: \quad 0 = \frac{\partial \Pi_{at}}{\partial p_{jt}} &= \sum_{k \in \mathcal{J}_{ft}} \frac{\partial}{\partial p_{jt}} \Big(p_{kt}\, s_{kt}(\cdot) - C_{kt}(\cdot) \Big) \; + \sum_{l \notin \mathcal{J}_{ft}} \frac{\partial}{\partial p_{jt}} \phi p_{lt}\, s_{lt}(\cdot) \\
&= s_{jt}(\cdot) + \sum_{k \in \mathcal{J}_{ft}} \Big((p_{kt} - c_{kt}) \frac{\partial s_{kt}(\cdot)}{\partial p_{jt}} \Big) + \sum_{l \notin \mathcal{J}_{ft}} \big(\phi p_{lt} \frac{\partial s_{lt}(\cdot)}{\partial p_{jt}} \big) \qquad (2.21) \\
\implies \quad s_{jt}(\cdot) &= - \sum_{k \in \mathcal{J}_{ft}} \Big(p_{kt} - c_{kt}(\cdot) \Big) \frac{\partial s_{kt}(\cdot)}{\partial p_{jt}} - \; \phi \sum_{l \notin \mathcal{J}_{ft}} p_{lt} \frac{\partial s_{lt}(\cdot)}{\partial p_{jt}} \\
&= \sum_{k \in \mathcal{J}_{ft}} \Big(p_{kt} - c_{kt}(\cdot) \Big) \Big(- \tfrac{\partial s_{kt}(\cdot)}{\partial p_{jt}} \Big) + \; \phi \sum_{l \notin \mathcal{J}_{ft}} p_{lt} \Big(- \tfrac{\partial s_{lt}(\cdot)}{\partial p_{jt}} \Big).
\end{aligned}
$$

$$
\begin{aligned}
\text{(b) If } j \notin \mathcal{J}_f: \quad 0 = \frac{\partial \Pi_{f't}}{\partial p_{jt}} &= \sum_{l \in \mathcal{J}_{f't}} \frac{\partial}{\partial p_{jt}} \Big((1-\phi)\, p_{lt}\, s_{lt}(\cdot) - C_{lt}(\cdot) \Big) \\
&= (1-\phi)\, s_{jt}(\cdot) + \sum_{l \in \mathcal{J}_{f't}} \Big((1-\phi)\, p_{lt} - c_{lt}(\cdot) \Big) \frac{\partial s_{lt}(\cdot)}{\partial p_{jt}}
\end{aligned}
$$

$$
\implies \quad (1-\phi)\, s_{jt}(\cdot) = \sum_{l \in \mathcal{J}_{f't}} \Big((1-\phi)\, p_{lt} - c_{lt}(\cdot) \Big) \big(- \frac{\partial s_{lt}(\cdot)}{\partial p_{jt}} \big). \qquad (2.22)
$$

In matrix form, let $\mathbf{s}_t, \mathbf{p}_t, \mathbf{c}_t, \mathbf{k}_t$ denote the vectors of shares, prices, costs and quality at time t. Define the following ownership matrices (see Appendix 2.A.2 for an example with three firms and six products):

- $\mathbf{H}_t^{all}$: each row corresponds to a product and indicates (with 1s and 0s) which products are owned by the same firm that owns the product in that row.

- $\mathbf{H}_t^{(f)}$: the submatrix of $\mathbf{H}_t^{all}$ that retains the original rows for products owned by firm f, while setting all other rows to zero. This captures the internal ownership structure of firm f.

- $\mathbf{H}_t^{(-f)}$: a matrix obtained from $\mathbf{H}_t^{all}$ by, for rows corresponding to products owned by firm f, replacing 0s with 1s and 1s with 0s, and setting all rows for products not owned by firm f to zero. This indicates, from firm f's perspective, the ownership of products held by the rival firms $f' \neq f$.

- $\mathbf{H}_t^{-f}$: a matrix with rows corresponding to products owned by firm f set to zero, and rows for products not owned by firm f equal to a vector of ones over all products not owned by f. This represents the ownership structure among the non-acquired firms, excluding products owned by f.

Consequently define:

- $\boldsymbol{\Delta}_t^p = \mathbf{H}_t^{all} \circ \left(-\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{p}_t}\right)$
- $\boldsymbol{\Delta}_t^{(f)p} = \mathbf{H}_t^{(f)} \circ \left(-\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{p}_t}\right)$
- $\boldsymbol{\Delta}_t^{(-f)p} = \mathbf{H}_t^{(-f)} \circ \left(-\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{p}_t}\right)$
- $\boldsymbol{\Delta}_t^{-fp} = \mathbf{H}_t^{-f} \circ \mathbf{H}_t^{all} \circ \left(-\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{p}_t}\right)$

We have seen earlier that the market share vector $\mathbf{s}_t$ depends on the variables $\mathbf{p}_t$, $\mathbf{g}_t$, and $\mathbf{d}_t$, and similarly, the $\boldsymbol{\Delta}_t$ matrices are functions of these variables as well. However, to simplify the notation below, we omit this explicit functional dependence. Assuming that firm f is the acquired firm, we express the relevant first-order conditions with respect to prices in a more compact form.

$$\begin{aligned}
\mathbf{s}_t &= \boldsymbol{\Delta}_t^{(f)p}\left(\mathbf{p}_t - \mathbf{c}_t\right) + \phi\boldsymbol{\Delta}_t^{(-f)p}\mathbf{p}_t + (1-\phi)\Delta_t^{-fp}\mathbf{p}_t - \boldsymbol{\Delta}_t^{-fp}\mathbf{c}_t \\
&= \left(\boldsymbol{\Delta}_t^{(f)p} + \phi\boldsymbol{\Delta}_t^{(-f)p} + (1-\phi)\boldsymbol{\Delta}_t^{-fp}\right)\mathbf{p}_t - \left(\boldsymbol{\Delta}_t^{(f)p} + \boldsymbol{\Delta}_t^{-fp}\right)\mathbf{c}_t. \qquad (2.23)
\end{aligned}$$

Define

$$\begin{aligned}
\mathbf{A}_p &\equiv \boldsymbol{\Delta}_t^{(f)p} + \phi\,\boldsymbol{\Delta}_t^{(-f)p} + (1-\phi)\,\boldsymbol{\Delta}_t^{-fp} \\
\mathbf{B}_p &\equiv \boldsymbol{\Delta}_t^{(f)p} + \boldsymbol{\Delta}_t^{-fp}, \qquad (2.24)
\end{aligned}$$

so that,

$$\mathbf{s}_t = \mathbf{A}_p\mathbf{p}_t - \mathbf{B}_p\mathbf{c}_t. \qquad (2.25)$$

Rearrange to solve for $\mathbf{c}_t$:

$$\begin{aligned}
\mathbf{B}_p\mathbf{c}_t &= \mathbf{A}_p\mathbf{p}_t + \mathbf{k}_t - \mathbf{s}_t \\
\mathbf{c}_t &= \mathbf{B}_p^{-1}\left(\mathbf{A}_p\mathbf{p}_t - \mathbf{s}_t\right). \qquad (2.26)
\end{aligned}$$

Once marginal costs are recovered, the implied markup vector for the acquired firm is

$$\mathbf{p}_t - \mathbf{c}_t = \mathbf{p}_t - \mathbf{B}_p^{-1}\left(\mathbf{A}_p\mathbf{p}_t - \mathbf{s}_t\right).$$

and the corresponding markup for the non-acquired firms is

$$(1-\phi)\mathbf{p}_t - \mathbf{c}_t = (1-\phi)\mathbf{p}_t - \mathbf{B}_p^{-1}\left(\mathbf{A}_p\mathbf{p}_t - \mathbf{s}_t\right). \tag{2.27}$$

Quality FOCs

$$\begin{aligned}
\text{(a) If } j \in \mathcal{J}_f: \quad 0 = \frac{\partial \Pi_{at}}{\partial g_{jt}} &= \sum_{k\in\mathcal{J}_{ft}} \frac{\partial}{\partial g_{jt}}\Big(p_{kt}s_{kt}(\cdot) - C_{kt}(\cdot)\Big) + \sum_{l\notin\mathcal{J}_{ft}} \frac{\partial}{\partial g_{jt}}\Big(\phi p_{lt}\, s_{lt}(\cdot)\Big) \\
&= -\frac{\partial C_{jt}(\cdot)}{\partial g_{jt}} + \sum_{k\in\mathcal{J}_{ft}} \Big(p_{kt} - c_{kt}(\cdot)\Big)\frac{\partial s_{kt}(\cdot)}{\partial g_{jt}} + \phi \sum_{l\notin\mathcal{J}_{ft}} p_{lt}\frac{\partial s_{lt}(\cdot)}{\partial g_{jt}} \\
\implies \quad \frac{\partial C_{jt}(\cdot)}{\partial g_{jt}} &= \sum_{k\in\mathcal{J}_{ft}} \Big(p_{kt} - c_{kt}(\cdot)\Big)\frac{\partial s_{kt}(\cdot)}{\partial g_{jt}} + \phi \sum_{l\notin\mathcal{J}_{ft}} p_{lt}\frac{\partial s_{lt}(\cdot)}{\partial g_{jt}} \\
\text{(b) If } j \notin \mathcal{J}_f: \quad 0 = \frac{\partial \Pi_{f't}}{\partial g_{jt}} &= \sum_{l\in\mathcal{J}_{f't}} \frac{\partial}{\partial g_{jt}}\Big((1-\phi)\, p_{lt}\, s_{lt}(\cdot) - C_{lt}(\cdot)\Big) \\
0 &= -\frac{\partial C_{jt}(\cdot)}{\partial g_{jt}} + \sum_{l\in\mathcal{J}_{f't}} \Big((1-\phi)p_{lt} - c_{lt}(\cdot)\Big)\frac{\partial s_{lt}(\cdot)}{\partial g_{jt}} \\
\implies \quad \frac{\partial C_{jt}(\cdot)}{\partial g_{jt}} &= \sum_{l\in\mathcal{J}_{f't}} \Big((1-\phi)p_{lt} - c_{lt}(\cdot)\Big)\frac{\partial s_{lt}(\cdot)}{\partial g_{jt}}
\end{aligned} \tag{2.28}$$

Matrix Form:

Let $\mathbf{g}_t$ denote the vector of qualities, and define for quality the analogous matrices:

- $\boldsymbol{\Delta}_t^{g} = \mathbf{H}_t^{all} \circ \left(\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{g}_j}\right)$
- $\boldsymbol{\Delta}_t^{(f)g} = \mathbf{H}_t^{(f)} \circ \left(\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{g}_t}\right)$
- $\boldsymbol{\Delta}_t^{(-f)g} = \mathbf{H}_t^{(-f)} \circ \left(\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{g}_t}\right)$
- $\boldsymbol{\Delta}_t^{-fg} = \mathbf{H}_t^{-f} \circ \mathbf{H}_t^{all} \circ \left(\frac{\partial \mathbf{s}_t(\cdot)}{\partial \mathbf{g}_t}\right)$

Then, the post-acquisition quality FOCs for all products can be written as:

$$\begin{aligned}
\mathbf{c}_t^g &= \boldsymbol{\Delta}_t^{(f)g}(\mathbf{p}_t - \mathbf{c}_t) + \phi\,\boldsymbol{\Delta}_t^{(-f)g}\,\mathbf{p}_t + (1-\phi)\,\boldsymbol{\Delta}_t^{-fg}\,\mathbf{p}_t - \boldsymbol{\Delta}_t^{-fg}\,\mathbf{c}_t \\
&= \Big(\boldsymbol{\Delta}_t^{(f)g} + \phi\,\boldsymbol{\Delta}_t^{(-f)g} + (1-\phi)\,\boldsymbol{\Delta}_t^{-fg}\Big)\,\mathbf{p}_t - \Big(\boldsymbol{\Delta}_t^{(f)g} + \boldsymbol{\Delta}_t^{-fg}\Big)\,\mathbf{c}_t,
\end{aligned} \tag{2.29}$$

where $\mathbf{c}_t^g$ is the vector of direct marginal quality costs, i.e., $[\,\partial C_{jt}/\partial g_{jt}\,]$ for all j.

Define

$$\begin{aligned}\mathbf{A}^g &\equiv \mathbf{\Delta}_t^{(f)g} + \phi\, \mathbf{\Delta}_t^{(-f)g} + (1-\phi)\, \mathbf{\Delta}_t^{-fg} \\ \mathbf{B}^g &\equiv \mathbf{\Delta}_t^{(f)g} + \mathbf{\Delta}_t^{-fg},\end{aligned} \tag{2.30}$$

so that

$$\mathbf{c}(\cdot)_t^g = \mathbf{A}^g \mathbf{p}_t - \mathbf{B}^g \mathbf{c}_t. \tag{2.31}$$

Unlike the standard Berry–Levinsohn–Pakes (BLP) framework, where marginal cost (of production and distribution) is treated as constant and exogenous, here the vector of marginal quality costs, $\boldsymbol{c}(\cdot)_t^g$, is modelled as a function of the market share functions and hence depends on $\boldsymbol{p}_t$, $\boldsymbol{g}_t$, and $\boldsymbol{d}_t$.

This generalization allows firms' cost structures to vary endogenously with quality, making quality a strategic decision variable alongside price. In this sense, $\mathbf{c}(\cdot)_t^g$ captures how costly it is for each firm to raise perceived or technological quality by one unit, rather than representing the marginal cost of production – how costly it is for each firm to increase the quantity by one unit.

In our empirical model, following Crawford et al. (2019), we specify an exponential functional form to characterize the relationship between product quality and production and distribution cost, so the total cost per capita is:

$$C_{jt}(\cdot) = \exp(c_{0jt} + c_{1jt} g_{jt}) s_{jt}(\mathbf{p}_t, \mathbf{g}_t, \mathbf{d}_t), \tag{2.32}$$

indicating that the marginal costs for production and distribution, and for quality, are respectively:

$$\begin{aligned}\partial C_{jt}/\partial s_{jt} &\equiv c_{jt}(\cdot) &&= \exp(c_{0jt} + c_{1jt} g_{jt}) \\ \partial C_{jt}/\partial g_{jt} &\equiv c_{jt}^g(\cdot) &&= \exp(c_{0jt} + c_{1jt} g_{jt}) c_{1jt} s_{jt}(\cdot) \\ dC_{jt}/dg_{jt} &\equiv c_{jt}(\cdot)\frac{\partial s_{jt}}{\partial g_{jt}} + c_{jt}^g(\cdot) &&= \exp(c_{0jt} + c_{1jt} g_{jt})\Big(c_{1jt} s_{jt}(\cdot) + \frac{\partial s_{jt}}{\partial g_{jt}}\Big).\end{aligned} \tag{2.33}$$

This specification allows for spillovers between quality and the processes of production and distribution in both directions.

2.3.3 Econometric specification

Our objective is to estimate both the consumer demand and firm supply parameters in a differentiated product market. This market is characterized by the endogenous determination of both prices and product qualities. The approach integrates structural econometric methods inspired by Berry et al. (1995), Nevo (2001), and Crawford et al. (2019).

2.3.3.1 Demand side

In each market t, consumer i chooses among J_t products and an outside option ($j = 0$) with utility normalized to zero. Indirect utility is

$$u_{ijt} \;=\; \delta_{jt} \;+\; \mu_{ijt} \;+\; \varepsilon_{ijt}, \qquad \varepsilon_{ijt} \sim \text{i.i.d. Type I extreme value,} \tag{2.34}$$

where δ_{jt} is the mean utility common to all consumers, μ_{ijt} captures taste heterogeneity, and ε_{ijt} is an idiosyncratic error.

Mean utility.

$$\delta_{jt} \;=\; \mathbf{x}_{jt}\boldsymbol{\beta} \;-\; \alpha\, p_{jt} \;+\; \tau\, g_{jt} \;+\; \xi_{jt}, \tag{2.35}$$

with $\mathbf{x}_{jt}$ observed characteristics,[17] p_{jt} price, g_{jt} perceived quality (e.g., ratings, possibly platform-influenced – due to lack of observability it is inclusive also of d_{jt}), and ξ_{jt} an unobserved demand shock known to firms. The linear demand parameters are $\boldsymbol{\theta}_1 \equiv \boldsymbol{\beta}$.

Heterogeneous tastes.

$$\mu_{ijt} \;=\; \alpha_i\, p_{jt} \;+\; \tau_i\, g_{jt} \;+\; \pi_p y_i p_{jt} \;+\; \pi_g y_i g_{jt} \;+\; \mathbf{v}_{jt}\boldsymbol{\Sigma}\boldsymbol{\nu_i}, \qquad \boldsymbol{\nu_i} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}), \tag{2.36}$$

where (α_i, τ_i) denote individual-specific sensitivities to price and perceived quality, y_i individual income, and $\boldsymbol{\Sigma}$ (diagonal) scales random tastes on additional product

[17] The vector of observed product characteristics $\mathbf{x}_{jt}$ is treated as exogenous because it reflects decisions taken before the product is introduced into the market. This can be interpreted as a stage 0, which precedes the platform's choice of quality $\mathbf{d}_t$. In principle, these observed characteristics should be time-invariant. We have nevertheless added a subscript t to allow for the inclusion of one time-varying observable characteristic, namely the observed age of the product. Notice also that, in every period, products enter and exit the market.

attributes $\mathbf{v}_{jt} \subseteq \mathbf{x}_{jt}$. In our empirical analysis $\mathbf{v}_{jt}$ is set to zero to focus on key heterogeneity associated with price and quality.

The nonlinear demand parameters are

$$\boldsymbol{\theta}_2 \equiv (\alpha, \tau, \sigma_p, \sigma_g, \pi_p, \pi_g, \boldsymbol{\Sigma}) \text{ simplified to } (\alpha, \tau, \sigma_p, \sigma_g, \pi_p, \pi_g).$$

Shares and inversion. Aggregating over the iid heterogeneity draws,

$$s_{jt}(\boldsymbol{\delta}_t(\boldsymbol{\theta}_2), \boldsymbol{\theta}_2) \;=\; \int \frac{\exp(\delta_{jt} + \mu_{ijt})}{1 + \sum_{k=1}^{J_t} \exp(\delta_{kt} + \mu_{ikt})}\, f(\alpha_i, \tau_i, y_i)\, d\alpha_i d\tau_i dy_i, \tag{2.37}$$

with the outside option in the denominator (+1). Given observed shares s_{jt} and a candidate $\boldsymbol{\theta}_2$, mean utilities are recovered by the Berry's fixed-point:

$$\boldsymbol{\delta}_t \;=\; \Delta^{-1}\big(\boldsymbol{s}_t; \boldsymbol{\theta}_2\big). \tag{2.38}$$

After inversion, the estimating mean utility equation is

$$\delta_{jt} \;=\; \mathbf{x}_{jt}\boldsymbol{\beta} \;-\; \alpha\, p_{jt} \;+\; \tau\, g_{jt} \;+\; \xi_{jt}, \tag{2.39}$$

Projection residuals (demand). The structural mean utility in Equation (2.39) can be specified in terms of the residual component:

$$\xi_{jt} \;=\; \delta_{jt} \;-\; \big(\mathbf{x}_{jt}\boldsymbol{\beta} \;-\; \alpha\, p_{jt} \;+\; \tau\, g_{jt}\big), \tag{2.40}$$

where ξ_{jt} is the *projection residual* from the demand equation. It captures unobserved demand shocks—such as brand reputation, latent product appeal, or unmeasured design attributes-that are known to firms but not observed by the econometrician.

In vector notation, the equation of interest is:

$$\boldsymbol{\xi}_t \;=\; \boldsymbol{\delta}_t \;-\; \big(\mathbf{X}_t\boldsymbol{\beta} \;-\; \alpha\, \mathbf{p}_t \;+\; \tau\, \mathbf{g}_t\big). \tag{2.41}$$

2.3.3.2 Supply side

The supply block follows the firm– and platform–optimization described in Section 2.3.2. We summarize the equilibrium conditions in compact matrix form and map them to estimable objects.

(i) Marginal cost of production and distribution Under multi-product Nash pricing with platform commission ϕ, following Equation 2.16, the marginal cost of production and distribution can be formulated as:

$$\boldsymbol{c}_t(\mathbf{g}_t) \;=\; (1-\phi)\Big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\mathbf{s}_t\Big). \tag{2.42}$$

Given the functional form of the marginal cost specified in Equation 2.33, Equation 2.42 can be reformulated as:

$$\mathbf{exp}(\boldsymbol{c}_{0t} + \boldsymbol{c}_{1t} \circ \boldsymbol{g}_t) \;=\; (1-\phi)\Big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\mathbf{s}_t\Big). \tag{2.43}$$

(ii) Marginal cost of quality By substituting Equation 2.16 into Equation 2.19 and rearranging the firms' first-order conditions with respect to $c_{jt}^g(\cdot)$, the marginal costs of quality can be expressed as:

$$\boldsymbol{c}_t^g(\mathbf{g}_t) \;\equiv\; (1-\phi)\boldsymbol{\Delta}_t^g(\cdot)\boldsymbol{\Delta}_t^p(\cdot)^{-1}\mathbf{s}_t \tag{2.44}$$

Given the functional form of the marginal cost specified in Equation 2.33, Equation 2.44 can be rewritten as:

$$\begin{aligned} \boldsymbol{c}_{1t} \circ \mathbf{exp}(\boldsymbol{c}_{0t} + \boldsymbol{c}_{1t} \circ \boldsymbol{g}_t) \circ \mathbf{s}_t \;&=\; (1-\phi)\boldsymbol{\Delta}_t^g(\cdot)\boldsymbol{\Delta}_t^p(\cdot)^{-1}\mathbf{s}_t \;\text{ or} \\ \boldsymbol{c}_{1t} \circ \mathbf{exp}(\boldsymbol{c}_{0t} + \boldsymbol{c}_{1t} \circ \boldsymbol{g}_t) \;&=\; (1-\phi)\boldsymbol{\Delta}_t^g(\cdot)\boldsymbol{\Delta}_t^p(\cdot)^{-1}\mathbf{1}_t. \end{aligned} \tag{2.45}$$

Cost functional form and projections. We project the level and slope of the marginal cost of production onto their corresponding shifters:

$$c_{0jt} = \mathbf{w}_{jt}\boldsymbol{\gamma} + \omega_{jt}, \qquad c_{1jt} = \mathbf{u}_{jt}\boldsymbol{\eta} + \upsilon_{jt}. \tag{2.46}$$

Divide (2.45) element-wise by (2.43) to obtain

$$\boldsymbol{c}_{1t} \;=\; \Big(\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t\Big) \;\oslash\; \Big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\Big), \tag{2.47}$$

that is, for each product j, $c_{1jt} = \big[\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t\big]_j \big/ \big[\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\big]_j$.
Substitute (2.47) into (2.43) and take logs element-wise:

$$\boldsymbol{c}_{0t} \;=\; \log\Big((1-\phi)\big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\big)\Big) \;-\; \Big((\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t) \;\oslash\; (\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t) \circ \boldsymbol{g}_t\Big), \tag{2.48}$$

i.e., for each j,

$$c_{0jt} = \log\Big((1-\phi)\big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\big)\Big)_j - \Big((\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t) \oslash (\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t) \circ \boldsymbol{g}_t\Big)_j.$$

Projection residuals. Considering the projections of c_{0jt} and c_{1jt} shown in Equation 2.46, the residual vectors are given by:

$$\begin{aligned}
\boldsymbol{\omega}_t &= \log\Big((1-\phi)\big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\big)\Big) - \bigg(\Big(\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t\Big) \;\oslash\; \Big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\Big)\bigg) \circ \boldsymbol{g}_t - \\
&\quad - \;\mathbf{W}_t\,\boldsymbol{\gamma} \\
\boldsymbol{\upsilon}_t &= \Big(\boldsymbol{\Delta}_t^g(\cdot)\,\boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\mathbf{1}_t\Big) \;\oslash\; \Big(\boldsymbol{p}_t - \boldsymbol{\Delta}_t^p(\cdot)^{-1}\,\boldsymbol{s}_t\Big) \;-\; \mathbf{U}_t\,\boldsymbol{\eta}.
\end{aligned} \tag{2.49}$$

Proxying the unknown platform cost One component of the pre-acquisition platform profit function in Equation 2.13 is the cost of platform-added quality, denoted by $K(d)$. Neither $K(d)$ nor the quality component d are observed in our data. We propose approximating the platform cost function using the platform's marginal revenue calculated at optimal seller qualities and prices.

Due to the unobservability of d, we compute the marginal revenue with respect to seller quality, following the spirit of Equation 2.1, where the effects of g and d are symmetric.

Therefore, the per-user seller (platform) marginal revenue of quality is given by

$$\phi\tau\mathbf{p}_t \circ \frac{\partial \mathbf{s}_t}{\partial \mathbf{g}_t}. \tag{2.50}$$

Assuming equality between marginal revenue and marginal cost, and further assuming constant marginal cost, we proxy the cost function as

$$\mathbf{k}_t(\mathbf{d}_t) = \phi\tau\mathbf{p}_t \circ \frac{\partial \mathbf{s}_t}{\partial \mathbf{g}_t} \circ \mathbf{d}_t. \tag{2.51}$$

Pre-acquisition one can think of a simplifying assumption $\mathbf{k}_t(\mathbf{d}_t = \mathbf{0}) = \mathbf{0}$. After the acquisition $\mathbf{k}_t(\Delta\mathbf{d}_t) = \phi\tau\mathbf{p}_t \circ \frac{\partial \mathbf{s}_t}{\partial \mathbf{g}_t} \circ \Delta\mathbf{d}_t$. This is the proxied change in platform cost after a change in quality, Δd.

2.3.4 Estimation approach

By stacking the systems of Equations 2.41 and 2.49 over time, we construct the system of simultaneous equations estimated using a non-linear GMM approach. This method addresses the endogenicity present in some of the variables. The number of equations (observations) is $3\sum_{t=1}^{T} J_t$.

Orthogonality conditions. Under the standard exogeneity assumptions, the structural residuals are mean independent of the corresponding observed characteristics. Since we are dealing with a system of simultaneous equations, the restriction required for identification is that

$$\mathbb{E}\big[\xi_{jt},\, \omega_{jt},\, \upsilon_{jt} \big|\, \mathbf{x}_{jt},\, \mathbf{w}_{jt}\,, \mathbf{u}_{jt}\big] = \mathbf{0}. \tag{2.52}$$

Indeed, some of the variables in $\mathbf{x}_{jt}$, $\mathbf{w}_{jt}$, and $\mathbf{u}_{jt}$ may overlap. It is convenient to rewrite the conditional moment conditions, which determine the functional forms of $\mathbf{x}_{jt}$, $\mathbf{w}_{jt}$, and $\mathbf{u}_{jt}$, in terms of instruments. Hence, we define $\mathbf{z}_{jt}^{(\xi)}$, $\mathbf{z}_{jt}^{(\omega)}$, and $\mathbf{z}_{jt}^{(\upsilon)}$ as the instruments associated with the first, second, and third set of equations, respectively.

These conditional moment restrictions imply the unconditional orthogonality conditions used in the stacked GMM estimation:

$$\mathbb{E}\big(\mathbf{z}_{jt}^{\xi}\, \xi_{jt}\big) = \mathbf{0}, \qquad \mathbb{E}\big(\mathbf{z}_{jt}^{\omega}\, \omega_{jt}\big) = \mathbf{0}, \qquad \mathbb{E}\big(\mathbf{z}_{jt}^{\upsilon}\, \upsilon_{jt}\big]) = \mathbf{0}. \tag{2.53}$$

Stacked moments and GMM criterion. For each market t with product set $\mathcal{J}_t$, define the stacked sample moments as

$$g_T(\boldsymbol{\theta}) \;=\; \frac{1}{T}\sum_{t=1}^{T}\frac{1}{J_t}\sum_{j\in\mathcal{J}_t}\begin{bmatrix}\mathbf{z}_{jt}^{\xi}\xi_{jt}(\boldsymbol{\theta})\\ \mathbf{z}_{jt}^{\xi}\omega_{jt}(\boldsymbol{\theta})\\ \mathbf{z}_{jt}^{\xi}\upsilon_{jt}(\boldsymbol{\theta})\end{bmatrix}, \tag{2.54}$$

where each block corresponds to one of the three projection equations on the demand side and the two supply sides. The linearity of the projection equations implies that $g_T(\boldsymbol{\theta})$ is linear in the coefficients $\boldsymbol{\theta}_1$ conditional on the nonlinear parameters $\boldsymbol{\theta}_2$, as in the standard BLP formulation.

The estimator solves the non-linear GMM problem

$$\arg\min_{\boldsymbol{\theta}}\; Q_T(\boldsymbol{\theta}) \;\equiv\; g_T(\boldsymbol{\theta})^{\top}\, W_T\, g_T(\boldsymbol{\theta}), \tag{2.55}$$

where W_T is a positive-definite weighting matrix of dimension the total number of instruments.

2.3.5 Identification Strategy

To identify the parameters, we employ several sets of instruments designed to mitigate potential endogenicity in price and quality. Our empirical strategy relies on cost-based instruments, BLP-type instruments, and differentiation-based instruments, each targeting specific sources of endogeneity.

First, following the approach of Berry et al. (1995), we construct BLP-type instruments based on product characteristics. These include both own-product characteristics, which account for how a product's specific features may influence its production cost or perceived value, and the characteristics of competing products, such as the average attributes of other firms' offerings. The latter captures how rivals' features affect a product's relative positioning in the market, thereby influencing firms' pricing and markup decisions. Together, these instruments help disentangle price variation driven by supply-side factors from that driven by demand-side preferences.

Second, drawing on Amit Gandhi and Jean-Francois Houde (2019), we construct excluded differentiation instruments that exploit observable differences across products. These instruments measure how each product's characteristics compare with those of both rival and non-rival goods. By capturing exogenous variation in product positioning, they reduce biases arising from unobserved demand shocks and further improve identification of structural parameters.

In our analysis, we also distinguish between horizontal and vertical product characteristics. Horizontal characteristics describe variations among products that cannot be ranked in terms of product quality,[18] such as physical dimensions including width, height, or weight, indicating differentiation based on taste or fit rather than superiority. Vertical characteristics, by contrast, reflect objective measures of quality or performance, such as battery life (in hours), suction power (measured in Pascals), or product age, all of which can be ranked according to desirability.

Overall, the integrated demand and supply framework enables us to capture the complex interplay between endogenous pricing and quality decisions in differentiated product markets.

2.4 Data and empirical results

2.4.1 Data

This study focuses on the Robotic vacuum cleaner industry, and specifically on iRobot, the largest producer of robotic vacuums worldwide. Table 2.1 illustrates that iRobot has been involved in a series of mergers and acquisitions. The recent merger agreement between Amazon and iRobot is valued at 1.7 billion dollars, which highlights the value of the robotic vacuum cleaner industry. However, because there hasn't been a horizontal merger in this industry, we use merger simulation to simulate the horizontal merger scenario.

2.4.1.1 Data sources

We collect data from four different sources (1) Keepa.com, (2) Amazon.com (3) AMZScout and (4) Ychart.com.

[18] This quality component is determined at the earlier stage 0 choice of product characteristics and therefore does not form part of g.

Measurement of quality Consumers rely on product ratings, comments, and reviews to determine a product's credibility and whether or not to purchase it. Therefore, we use the product's historical average rating on Keepa.com as a proxy for the quality of the product.

Product data From Keepa.com (hereafter, Keepa), we gather product-level information. Since 2011, Keepa, a private online market intelligence company, has been scraping Amazon for Amazon buyers and sellers. Keepa gives each product two different kinds of data: high-frequency time series data at the product level and product-specific descriptive data.

Estimating sales quantities Keepa only provides information on sales rank. We convert daily sales ranks to sales amounts using estimates from AMZScout[19], the market intelligence leader for Amazon sellers.

Market definition To define the overall market, we utilize historical monthly data on US households and income percentages provided by YCharts.[20] Based on this data, our analysis assumes that 14% of American households-consistent with the figures reported in the Smart Robot Vacuum Cleaners Statistics of 2023-[21] will purchase a robotic vacuum cleaner.

2.4.1.2 Descriptive statistics

To ground our empirical analysis, we begin with a detailed summary of the key variables used in the estimation, focusing on both product-level and firm-level characteristics.

Table 2.4 provides a comprehensive overview of the primary variables in our dataset, including measures of product performance such as average ratings and prices, as well as physical and technical specifications like dimensions, weight, and suction power. The table indicates substantial heterogeneity across products, as reflected in the wide ranges of the variables. For instance, prices span from $29.99 to $3,099, highlighting significant market segmentation. Similarly, suction power varies from 600 Pa to 6,800 Pa, indicating a diverse range of product quality and performance.

The average rating of 4.08 suggests that most products are favourably reviewed,

[19] https://amzscout.net/sales-estimator/
[20] https://ycharts.com/indicators/us_households/chart/
[21] https://tinyurl.com/2bu356p5

with relatively low variation (standard deviation of 0.51). These summary statistics establish the foundation for the subsequent econometric analyses, which explore how these characteristics influence consumer choice and firm behaviour.

Table 2.4: Summary statistics of relevant variables

Variable	**N**	**Mean**	**Std. Dev.**	**Min**	**Max**
Average Rating	3,592	4.08	0.51	1.00	5.00
Price (in 100 $)	3,592	3.48	2.34	0.30	30.99
Market share (s)	3,592	3e-4	0.001	5.55e-08	0.01
Width (inches)	3,592	12.65	3.55	1.72	23.35
Height (inches)	3,592	6.01	5.82	0.92	33.38
Weight (lbs)	3,592	8.71	4.23	0.65	32.02
Battery Life (hours)	3,592	1.85	0.56	1.00	3.33
Suction Power (KPa)	3,592	2.20	1.17	0.60	6.800
Product Age (months)	3,592	66.38	32.69	15.00	155.00
FBA Fees ($)	3,592	8.34	2.63	2.41	22.32
HHINCOME (in 100k$)	6,900	0.82	0.54	0.10	2.30

Notes: This table reports descriptive statistics for robotic vacuum cleaners sold in the US market. *Average Rating* is based on user reviews; *Price* is reported in 100$ increments; *Market share* includes also the outside option of not-buying; *Width*, *Height*, and *Weight* capture physical dimensions; *Battery Life* is measured in hours of operation per full charge; *Suction Power* is measured in kilopascals (KPa); *Product Age* is the time since product launch (in months); *FBA Fees* denote Fulfillment by Amazon charges in US dollars. From the IPUMS data (`https://ipums.org/`), we randomly draw 100 individuals for each of the 69 markets (periods). To avoid issues with outliers, we restrict the income variable to its 5th and 95th percentile (10,000 and 230,419.6 United States Dollar (USD), respectively). We then rescale income by dividing it by 100,000. Sample sizes vary due to data availability.

Table 2.5: Number of firms and brands

Statistic	**T**	**Mean**	**Std. Dev.**	**Min**	**Max**
Number of Firms	69	13.62	5.15	5	23
Number of Brands	69	52.06	24.23	15	115

Notes: Observations are monthly and are considered individual markets. The table summarizes the variation in number of firms and brands over time.

Table 2.5 highlights the distribution of firms and brands across markets. On average, each market contains 14 firms and 52 products, with notable variation. The maximum of 115 products in a single market underscores the intense competition in some segments of the robotic vacuum cleaner industry. This diversity in market structure provides a rich context for examining the interaction between market characteristics and firm strategies.

2.4.2 Estimation results

Table 2.6 reports the joint estimation results from the random-coefficients Logit demand model and the two supply-side cost equations. As expected, accounting for endogenous pricing and quality substantially alters the demand estimates relative to OLS. In particular, the GMM specification yields a much larger price sensitivity: the mean price coefficient is nearly six times larger in magnitude than its OLS counterpart, consistent with the well-known downward bias of price coefficients when price endogenicity is ignored. Consumers also respond strongly and positively to perceived quality, proxied by product ratings. The estimated rating coefficient more than doubles relative to OLS, confirming that ratings are a first-order driver of demand in this market once simultaneity between quality, prices, and unobserved demand shocks is addressed.

Among objective product characteristics, suction power enters positively and significantly, indicating that consumers value measurable cleaning performance. Battery life, by contrast, switches sign relative to OLS and becomes positive but statistically imprecise in the GMM specification, suggesting that its apparent negative effect in OLS reflects correlation with unobserved product attributes—such as energy consumption or other design trade-offs-that enter consumers' utility but are not separately observed, rather than a direct dis-utility from longer battery life.

The random-coefficients estimates reveal substantial heterogeneity in preferences. The estimated standard deviations of the price and rating coefficients indicate meaningful dispersion in consumer sensitivity to both price and perceived quality. The income interaction terms suggest systematic variation across consumers: higher-income consumers exhibit greater price sensitivity and place relatively less weight on ratings, consistent with heterogeneous valuation patterns in this product category.

A mean price coefficient of -1.951 means that a \$100 increase in a vacuum cleaner's price multiplies the odds of choosing it over the outside option by $\exp(-1.951) \approx 0.142$, or roughly a $(0.142-1)\times 100 = 85.79\%$ reduction in those odds. Similarly, a coefficient

Table 2.6: Joint estimation results: Demand and cost-component equations

	Ordinary Least Squares (OLS)		GMM RCL	
Variable	**Estimate**	**SE**	**Estimate**	**SE**
Demand Equation				
Constant	-17.716	1.937	-25.243	2.898
Suction Power	1.003	0.394	2.290	0.255
Battery Life	-2.959	0.930	1.113	0.846
Height	-0.056	0.093	0.004	0.038
Price ($-\alpha$)	-0.348	0.050	-1.951	0.020
Rating (τ)	0.834	0.129	1.998	0.020
Price (σ_α)	-	-	0.472	0.011
Price (π_α)	-	-	-2.321	0.058
Rating (σ_τ)	-	-	0.333	0.007
Rating (π_τ)	-	-	-1.060	0.045
$\mathbf{c_0}$ Equation				
Constant	-	-	-24.164	3.303
Height	-	-	0.007	0.090
Weight	-	-	0.330	0.227
$\mathbf{c_1}$ Equation				
Constant	-	-	8.725	6.842
Suction Power	-	-	-0.199	0.192
Battery Life	-	-	-0.484	2.072
Model Statistics				
N	3592	-	3592	-
J-stat	-	-	2	-
N MC (<0)	-	-	20	-
Average PCM	-	-	0.261	-
Average $\mathbf{c_t}$	-	-	2.42	-
Average $\mathbf{c_t^g}$	-	-	0.0004	-

Notes: Columns (1)–(2) report OLS estimates of the demand equation. Columns (3)–(4) report joint non-linear GMM estimates from a random-coefficients logit model with endogenous price and rating. Random coefficients are allowed on price and rating. Demographic variables (including income) are scaled and demeaned at the market level. The c_0 and c_1 equations correspond to the level and slope of marginal costs. All specifications include Month $\times$ Firm fixed effects. Robust standard errors are reported.

on rating of 1.998 means that an increase in the ratings by one unit (an extreme value of two standard deviations) multiplies the odds of choosing that product over the outside option by $\exp(1.998) \approx 7.37$, or roughly a $(7.397 - 1) \times 100 = 637\%$ increase in those odds.

Since the income variable is demeaned, these results approximate the effects for an individual with average income. For the richest individual ($230K), the price effect becomes

$$\exp\Big(-1.951 - 2.321 \times (2.30 - 0.82)\Big) = \exp(-5.386) \approx 0.0046,$$

corresponding to a change in odds of

$$(0.0046 - 1) \times 100 \approx -99.54\%,$$

which indicates much greater price sensitivity among richer individuals.

By contrast, the price effect for the poorest individual is

$$\exp\Big(-1.951 - 2.321 \times (0.10 - 0.82)\Big) = \exp(-0.281) \approx 0.755,$$

or

$$(0.755 - 1) \times 100 \approx -24.5\%$$

change in odds.

The same logic applied to quality yields, for the richest individual,

$$\exp\Big(1.998 - 1.06 \times (2.30 - 0.82)\Big) = \exp(0.429) \approx 1.54,$$

i.e.

$$(1.54 - 1) \times 100 \approx 54\%$$

higher odds, suggesting substantially reduced sensitivity to quality relative to price at high incomes. For the poorest individual, the quality effect is

$$\exp\Big(1.998 - 1.06 \times (0.10 - 0.82)\Big) = \exp(2.761) \approx 15.8,$$

or

$$(15.8 - 1) \times 100 \approx 1{,}480\%$$

higher odds. These interpretations exclude the additional random heterogeneity captured by the standard deviations (0.472 for price and 0.333 for quality).

Richer individuals exhibit greater price sensitivity but substantially reduced responsiveness to quality improvements, while poorer individuals show the opposite: moderate price tolerance but extreme quality sensitivity. For online purchases of robotic vacuum cleaners, this pattern may be explained by income-driven search costs: affluent buyers, often equipped with premium models already, treat upgrades as marginal conveniences and scrutinize prices closely, whereas lower-income buyers prioritize durable quality gains as infrequent, high-stakes investments justifying higher relative tolerance for cost.

Turning to the supply side, the recovered marginal-cost parameters suggest that physical characteristics may act as cost shifters, although these effects are imprecisely estimated. Height and weight enter the marginal production cost equation with positive coefficients, consistent with higher engineering and logistics costs for larger devices, but neither effect is statistically significant. Estimates of the quality-cost slope indicate that the marginal cost of improving quality varies across performance attributes; however, these parameters are estimated with substantial uncertainty, and their signs should be interpreted cautiously.

The average marginal cost of production is estimated at \$242, while the marginal cost of an additional star rating is only about 4 cents -a substantially lower budget.

Overall, the joint demand–supply estimates point to strong consumer responsiveness to both prices and ratings, alongside meaningful cost–quality interactions on the supply side. These features are central to the analysis that follows. In particular, the combination of high demand elasticity with respect to perceived quality and non-trivial quality costs implies that platform-driven shifts in ratings can materially alter firms' optimal pricing and real quality choices. This provides the empirical foundation for the counterfactual simulations in Section 2.4.2.1, where exogenous rating advantages for iRobot are shown to reallocate market shares, profits, and surplus even in the absence of changes in underlying technological quality. Finally, Appendix 2.A.3 confirms the relevance of the IV, with first-stage

F-statistics well above conventional thresholds for both price and ratings.

2.4.2.1 Counterfactual analysis

We simulate the effect of an exogenous increase in iRobot's perceived rating by introducing a utility wedge d into its mean utility. This wedge captures a platform-driven boost to perceived quality and does not affect production technology or cost fundamentals (stage 0). In constructing the counterfactual, we hold the estimated cost primitives-namely the parameters governing marginal production costs and the marginal cost of firm-controlled quality-fixed at their pre-counterfactual values. Given these fixed cost primitives, we re-solve for the firms' joint pricing and quality equilibrium under the altered demand environment.

Consistent with the model's definition of quality, the variable g represents a firm's endogenous choice of overall product quality or appeal, encompassing both tangible attributes (such as design or durability) and intangible components (such as branding, customer support, or service quality). Observed technical characteristics—including suction power, battery life, and physical dimensions-are held fixed throughout the analysis. The counterfactual therefore isolates the impact of a platform-driven rating advantage by allowing only prices and firms' endogenous quality choices g to adjust in equilibrium.

We compute the new equilibrium for values of the platform-driven perceived-rating change $\Delta d \in \{0.10, 0.20, 0.30, 0.40, 0.50\}$. Throughout the analysis, reported changes in prices and endogenous quality are share-weighted across inside goods, so that products with negligible market shares do not receive disproportionate weight. Table 2.7 reports the resulting outcomes for Market 69 (the final period) and highlights three systematic mechanisms through which a perceived rating advantage reshapes market equilibrium.

First, the exogenous rating boost raises iRobot's perceived quality in consumers' utility, which shifts demand toward iRobot and increases the marginal payoff to further seller-side quality investment. In the model, g is chosen endogenously by the firm and captures intrinsic product quality improvements under the firm's control. Once the platform-driven rating boost expands demand, an additional increase in g becomes more profitable because it improves the attractiveness of a now larger and more valuable customer base. As a result, iRobot optimally responds not only by adjusting price, but also by increasing its endogenous quality choice g. Across counterfactuals, iRobot's price falls by roughly 29–41 percent, while its endogenous quality increases by about 20–23 percent. This pattern reflects the fact that the rating advantage induces substantial demand expansion, making aggressive price reductions profitable for the privileged product, while higher endogenous quality further amplifies the perceived-quality advantage in consumer utility.

Second, rival firms adjust despite not receiving the rating wedge directly. The perceived quality improvement for a prominent product expands demand for the category as a whole, leading to a sizeable contraction of the outside option share (around 8–14 percentage points). Rivals' market shares therefore rise in proportional terms, and their equilibrium responses differ from iRobot's: rivals raise prices modestly (about 0.5–6 percent) while also increasing endogenous quality (about 8–17 percent). Market expansion and reduced outside-option substitution allow rivals to increase markups even as iRobot prices fall.

Third, these firm-level adjustments translate into pronounced aggregate effects. Because demand reallocates strongly toward iRobot's sharply lower-priced products, the share-weighted average price across inside goods declines markedly (about 23–33 percent), despite small price increases among rival products. At the same time, share-weighted endogenous quality rises at the block level (about 15–20 percent), and consumer surplus increases substantially. In levels, average per-consumer surplus rises from approximately 0.01 in the baseline equilibrium to about 0.06 following the rating intervention, which corresponds to roughly \$1 to \$6 per consumer given that prices are scaled in units of \$100. The large percentage increase in consumer surplus therefore reflects a low baseline level combined with sizeable entry from the outside option, rather than large welfare gains for existing consumers.

Finally, profit impacts reflect both market expansion and the platform accounting. To avoid double counting, we compute Amazon's platform profit and iRobot's producer profit separately, treating the two entities as operating independently for accounting purposes, even though the counterfactual is motivated by a proposed acquisition. Platform profit is defined as commission revenue net of the cost of providing perceived quality, where the cost of the platform-driven rating contribution Δd is proxied using Equation 2.51. Under this accounting, iRobot's stand-alone profit rises strongly, rivals' profits also increase, and platform profit increases for smaller rating shocks but declines for larger interventions.

Taken together, the counterfactual analysis shows that platform-driven rating privileges operate through a joint channel of perceived quality, market expansion, and equilibrium reallocation of demand. Even modest algorithmic interventions—well within observed variation in online ratings-can therefore materially alter equilibrium prices, quality choices, participation, and surplus allocation in platform-mediated markets.

Table 2.7: Counterfactual rating privilege for iRobot (Market 69)

Panel A. Competitive responses: iRobot vs. others

Δ (stars)	Δp (%)	iR Δp (%)	Others Δp (%)	Δg (%)	iR Δg (%)	Others Δg (%)
0.10	-33.28	-40.97	+0.50	+15.38	+19.77	+8.50
0.20	-27.22	-33.75	+1.75	+17.40	+22.10	+10.04
0.30	-22.73	-29.05	+5.89	+18.47	+22.59	+13.00
0.40	-24.37	-31.69	+4.12	+17.86	+21.23	+14.71
0.50	-28.56	-34.17	+3.54	+19.69	+22.89	+16.81

Panel B. Market-level and platform outcomes

ΔOutside share (pp)	iR Δ share (%)	Others Δ share (%)	ΔCS (%)	$\Delta\pi_{\text{iRobot}}$ (%)	$\Delta\pi_{\text{Other}}$ (%)	$\Delta\pi_{\text{Platform}}$ (%)
-7.787	+803.50	+64.94	+391.36	+749.08	+100.67	+135.32
-8.271	+832.20	+82.05	+418.67	+734.66	+122.57	+96.79
-8.395	+840.45	+85.89	+426.03	+712.67	+127.66	+33.53
-9.675	+931.47	+121.81	+498.08	+841.89	+174.25	-19.77
-13.707	+1375.31	+138.38	+738.10	+1439.49	+194.45	-115.38

Notes: Δ (stars) denotes the exogenous increase in iRobot's perceived rating, measured in star units on the platform's 1–5 rating scale. Perceived ratings are capped at 5 stars. In the baseline equilibrium of Market 69, iRobot's average share is about 0.01, other inside products together account for roughly 0.013, and the outside option captures approximately $\approx$ 0.977 of demand. Panel A reports firm-level percentage changes for iRobot (iR) and its rivals following exogenous increases in perceived ratings. Panel B reports market-level changes in consumer surplus (CS), the outside option share, and profit changes for iRobot, rivals, and the platform. All outcomes are relative to the baseline equilibrium in Market 69. Market share changes are reported in proportional terms; changes in the outside option are in percentage points. CS is computed per consumer from the logit inclusive value. In levels, average CS rises from about 0.01 to 0.06. Since prices are scaled in units of $100, these correspond to roughly $1 and $6 per consumer. The large percentage increase therefore reflects a low baseline CS combined with substantial inflows from the outside option. Reported price and quality changes are share-weighted across inside goods.

2.5 Conclusion and discussion

This chapter shows that platform-controlled quality signals can be a quantitatively important channel of competitive harm in digital markets. Even modest improvements in a target firm's perceived quality, implemented through platform-mediated rating or review mechanisms, can generate large shifts in market share, profits, and equilibrium conduct. Importantly, these effects are not limited to consumer perception alone: they also alter firms' incentives to adjust prices and invest in seller-controlled product quality. The broader implication is that merger assessment in digital-platform settings should not focus only on prices or traditional foreclosure mechanisms. It should also examine how ownership may change the platform's incentives to steer visibility, ratings, and other information signals in ways that reshape competition. Substantively, our estimates imply that a vertically integrated platform can engineer large competitive advantages for an acquired brand without touching the physical product or posting higher prices: a small, algorithmically delivered rating advantage is sufficient to expand the integrated firm's demand, reallocate market share away from rivals, and raise both firm and platform profits. In other words, the merger's competitive bite operates primarily through perceived quality and demand expansion rather than through the price and foreclosure channels that conventional merger analysis is designed to detect.

This paper provides new evidence on the competitive effects of digital platform acquisitions by analysing Amazon's proposed purchase of iRobot through the lens of platform-driven quality contribution. By integrating a theoretical framework with structural empirical analysis, we demonstrate that platforms possess powerful non-price instruments to alter market outcomes following vertical integration. Our results show that even modest, algorithmic improvements to the acquired firm's ratings can yield substantial shifts in market share and profitability, conferring durable competitive advantages that extend well beyond traditional price effects.

These findings highlight the central role of consumer ratings as a strategic lever in digital markets, and illustrate how platform ownership can amplify quality-based competition to the benefit of the integrated firm. Our counterfactual simulations reveal that the ability to selectively enhance the perceived quality of an acquired brand can fundamentally reshape the allocation of surplus and competitive dynamics within the market. More broadly, the results suggest that non-price mechanisms are not peripheral in platform markets; they may be central to how market power is exercised after integration.

The implications for competition policy are clear: conventional merger analysis, with its predominant focus on prices and output, may fail to capture the full scope of welfare and market power distortions in platform-mediated environments. Our evidence suggests that antitrust authorities must pay greater attention to the potential for subtle, algorithmic quality contributions to entrench dominant positions and harm rivals, even in the absence of overt price manipulation or exclusion.

More broadly, this study contributes to the literature on digital competition by documenting a novel channel through which platforms can influence market structure and consumer welfare. In markets where product differentiation hinges critically on perceived quality, the intersection of platform power and quality signals creates new regulatory challenges. Future research should investigate the long-run effects of platform-driven quality contribution on innovation, entry, and consumer trust, as well as the design of effective policy tools for digital merger review.

It is worth concluding with a brief discussion of a striking real-world development that falls outside the scope of this article. On 14 December 2025, iRobot filed for bankruptcy. The bankruptcy of iRobot following the collapse of Amazon's proposed acquisition offers a revealing contrast to the counterfactual scenario analysed in this work. In the model, Amazon acquires iRobot and exploits platform-mediated quality signals to reshape demand, insulating the firm from price and quality competition despite unchanged production costs. By contrast, in reality iRobot remained independent, entered bankruptcy, and was ultimately acquired by its Chinese supplier into a market space dominated by low-cost competitors. Lacking access to platform-driven demand advantages, iRobot was forced to compete primarily on cost, where it faced a structural disadvantage. This comparison suggests that vertical integration with a platform may matter less for cost efficiencies than for sustaining perceived quality and market power through demand-side channels.

The European Commission's decision to challenge the merger was grounded in a rigorous technical assessment, and the analysis in this work is broadly consistent with that decision. Yet, as this episode illustrates, an independent decision taken under uncertainty can lead to outcomes that differ markedly from what informed policymakers and observers might have anticipated at the time. In retrospect, and especially when viewed through the lens of subsequent political commentary, there is a risk that such decisions are judged with the benefit of hindsight or framed in terms of political bias rather than the economic evidence available when they were taken. Comments from this Financial Times article provide a flavour of such bias `https://tinyurl.com/3wavkrud`: 'Regulators k*led the iRobot' `https://tinyurl.com/y9nfcvxa`; 'What again was the economic justification to block Amazon's bid to buy iRobot? The company clearly thought they could make a better product building on the Roomba base and would integrate it into their smart home product suite. Good on them. And Amazon was willing to accept the condition that they would not block Chinese competitors from their ecommerce marketplace. The EU killed the business plan. They instead forced a bankruptcy and now one the Chinese robot vacuum makers picks up some tech and a brand on the cheap. Great use of government power by the EU. Bravo!' https://tinyurl.com/mrx88dwt; 'Well done EU regulations, once again being Chinese puppets.' https://tinyurl.com/3t9jbfdz.

2.A Appendix

2.A.1 Derivations

2.A.1.1 Firms' price and quality best responses

Taking (d_1, d_2) as given, each firm solves

$$\max_{p_i, g_i} \ \pi_i = (1-\phi)\, p_i\, q_i \ - \ c\, q_i \ - \ \tfrac{1}{2} g_i^2, \quad q_i = \lambda_i + \beta(g_i + d_i) + \rho\, p_j - p_i,$$

for $i \neq j$. The first-order conditions yield

$$p_1^*(d_1,d_2) = \frac{-c\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^4 + d_1\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^3 - (\phi-1)\,((\rho\lambda_2+\lambda_1)\,\phi + c\,\rho^2 - \rho\lambda_2 - 3c - \lambda_1)\,\beta^2 + (d_1\,\rho^2 + d_2\rho - 2d_1)\,(\phi-1)\,\beta + (-\rho\lambda_2 - 2\lambda_1)\,\phi + (c+\lambda_2)\,\rho + 2c + 2\lambda_1}{((\phi-1)\,(\rho-1)\,\beta^2 + \rho - 2)\,(\phi-1)\,((\rho+1)\,(\phi-1)\,\beta^2 + \rho + 2)},$$

$$p_2^*(d_1,d_2) = \frac{-c\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^4 + d_2\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^3 - (\phi-1)\,((\lambda_1\rho+\lambda_2)\,\phi + c\,\rho^2 - \lambda_1\rho - 3c - \lambda_2)\,\beta^2 + (d_2\,\rho^2 + d_1\rho - 2d_2)\,(\phi-1)\,\beta + (-\lambda_1\rho - 2\lambda_2)\,\phi + (c+\lambda_1)\,\rho + 2c + 2\lambda_2}{((\phi-1)\,(\rho-1)\,\beta^2 + \rho - 2)\,(\phi-1)\,((\rho+1)\,(\phi-1)\,\beta^2 + \rho + 2)},$$

$$g_1^*(d_1,d_2) = -\frac{\left(d_1\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^3 + (c\,\rho^2 - \lambda_2\,(\phi-1)\,\rho - \lambda_1\phi - c + \lambda_1)\,(\phi-1)\,\beta^2 + (d_1\,\rho^2 + d_2\rho - 2d_1)\,(\phi-1)\,\beta + c\,\rho^2 + (-\lambda_2\phi + c + \lambda_2)\,\rho - 2\lambda_1\phi - 2c + 2\lambda_1\right)\beta}{((\phi-1)\,(\rho-1)\,\beta^2 + \rho - 2)\,((\rho+1)\,(\phi-1)\,\beta^2 + \rho + 2)},$$

$$g_2^*(d_1,d_2) = -\frac{\left(d_2\,(\rho-1)\,(\rho+1)\,(\phi-1)^2\,\beta^3 + (c\,\rho^2 - \lambda_1\,(\phi-1)\,\rho - \lambda_2\phi - c + \lambda_2)\,(\phi-1)\,\beta^2 + (d_2\,\rho^2 + d_1\rho - 2d_2)\,(\phi-1)\,\beta + c\,\rho^2 + (-\lambda_1\phi + c + \lambda_1)\,\rho - 2\lambda_2\phi - 2c + 2\lambda_2\right)\beta}{((\phi-1)\,(\rho-1)\,\beta^2 + \rho - 2)\,((\rho+1)\,(\phi-1)\,\beta^2 + \rho + 2)}$$

2.A.1.2 Optimal platform rating adjustment without acquisition

Substitute the above p_i^*, g_i^* into

$$\Pi_a(d_1,d_2) = \phi\big(p_1^* q_1^* + p_2^* q_2^*\big) - \tfrac{1}{2}(d_1^2 + d_2^2),$$

and solve $\partial \Pi_a / \partial di = 0$. The resulting optimal contributions are

$$d_1^* = -\frac{\phi\left(c\,(\phi-1)^3\,\beta^6 + 3\,(\phi-1)\left(-\frac{2\lambda_1\,\phi^2}{3} + \left(c\rho + \frac{2}{3}c + \frac{4}{3}\lambda_1\right)\phi - c\rho - \frac{4c}{3} - \frac{2\lambda_1}{3}\right)\beta^4 + ((-4\lambda_2\rho - 4\lambda_1)\,\phi^2 + (3c\,\rho^2 + (6c + 8\lambda_2)\,\rho + 4c + 12\lambda_1)\,\phi - 3c\,\rho^2 + (-8c - 4\lambda_2)\,\rho - 4c - 8\lambda_1)\,\beta^2 + (-2\lambda_1\,\rho^2 - 8\lambda_2\rho - 8\lambda_1)\,\phi + c\,\rho^3 + (4c + 2\lambda_1)\,\rho^2 + (4c + 8\lambda_2)\,\rho + 8\lambda_1\right)\beta}{\left((\phi-1)^2\,\beta^4 + ((2\rho+2)\,\phi - 2\rho - 4)\,\beta^2 + (\rho+2)^2\right)(\phi-1)\left((\phi-1)^2\,\beta^4 + ((-2\rho+2)\,\phi + 2\rho - 4)\,\beta^2 + (\rho-2)^2\right)},$$

$$d_2^* = -\frac{\phi\left(c\,(\phi-1)^3\,\beta^6 + 3\,(\phi-1)\left(-\frac{2\lambda_2\,\phi^2}{3} + \left(c\rho + \frac{2}{3}c + \frac{4}{3}\lambda_2\right)\phi - c\rho - \frac{4c}{3} - \frac{2\lambda_2}{3}\right)\beta^4 + ((-4\lambda_1\rho - 4\lambda_2)\,\phi^2 + (3c\,\rho^2 + (6c + 8\lambda_1)\,\rho + 4c + 12\lambda_2)\,\phi - 3c\,\rho^2 + (-8c - 4\lambda_1)\,\rho - 4c - 8\lambda_2)\,\beta^2 + (-2\lambda_2\,\rho^2 - 8\lambda_1\rho - 8\lambda_2)\,\phi + c\,\rho^3 + (4c + 2\lambda_2)\,\rho^2 + (4c + 8\lambda_1)\,\rho + 8\lambda_2\right)\beta}{\left((\phi-1)^2\,\beta^4 + ((2\rho+2)\,\phi - 2\rho - 4)\,\beta^2 + (\rho+2)^2\right)(\phi-1)\left((\phi-1)^2\,\beta^4 + ((-2\rho+2)\,\phi + 2\rho - 4)\,\beta^2 + (\rho-2)^2\right)}.$$

2.A.1.3 Optimal platform rating adjustment with acquisition

When the platform acquires firm 1, its objective becomes

$$\Pi_a(d_1,d_2) = (p_1^* - c)\, q_1^* \ - \ \tfrac{1}{2}\, g_1^{*2} \ + \ \phi\, p_2^*\, q_2^* \ - \ \tfrac{1}{2}(d_1^2 + d_2^2),$$

The first-order conditions

$$\frac{\partial \Pi_a}{\partial d_1} = 0, \quad \frac{\partial \Pi_a}{\partial d_2} = 0$$

yield closed-form solutions

$$d_1^* = -\frac{\begin{bmatrix} c(\rho-1)^2(\rho+1)^2(\phi-1)^3\beta^6 - \big(c(\phi+1)\rho^2 - (\phi-1)(2\lambda_2\phi + c - 2\lambda_2)\rho \\ \quad - 2\lambda_1\phi^2 + (2c+4\lambda_1)\phi - 4c - 2\lambda_1\big)(\rho-1)(\rho+1)(\phi-1)\beta^4 \\ + (\rho-1)\Big(c(\phi+1)\rho^3 + (c-4\lambda_2)\phi + 3c + 4\lambda_2)\rho^2 \\ \quad + [4\lambda_1\phi^2 + (-2c - 12\lambda_1 - 4\lambda_2)\phi + 4c + 8\lambda_1 + 4\lambda_2]\rho \\ \quad - 4(\phi-1)(-\lambda_1\phi + c + 2\lambda_1)\Big)\beta^2 \\ - c\rho^4 + (2\lambda_2\phi - 3c - 2\lambda_2)\rho^3 + 6\lambda_1(\phi-1)\rho^2 + 4c\rho - 8\lambda_1(\phi-1) \end{bmatrix}\phi\beta}{\begin{bmatrix} (\phi-1)^2(\rho+1)^2\beta^4 + 2(\rho+1)(-\rho+\phi-2)\beta^2 \\ + (\rho+2)^2 \end{bmatrix}(\phi-1)\begin{bmatrix} (\phi-1)^2(\rho-1)^2\beta^4 - 2(\rho-1)(\rho+\phi-2)\beta^2 \\ + (\rho-2)^2 \end{bmatrix}}$$

$$d_2^* = -\frac{\begin{bmatrix} c(\rho-1)^2(\rho+1)^2(\phi-1)^3\beta^6 - (\rho-1)(\rho+1)(\phi-1) \\ \quad \cdot \big(c(\phi+1)\rho^2 - (\phi-1)(2\lambda_1\phi + c - 2\lambda_1)\rho - 2\lambda_2\phi^2 + (2c+4\lambda_2)\phi - 4c - 2\lambda_2\big)\beta^4 \\ + (\rho-1)\Big(c(\phi+1)\rho^3 + [c-4\lambda_1)\phi + 3c + 4\lambda_1]\rho^2 \\ \quad + [4\lambda_2\phi^2 + (-2c - 4\lambda_1 - 12\lambda_2)\phi + 4c + 4\lambda_1 + 8\lambda_2]\rho \\ \quad - 4(\phi-1)(-\lambda_2\phi + c + 2\lambda_2)\Big)\beta^2 \\ - c\rho^4 + (2\lambda_1\phi - 3c - 2\lambda_1)\rho^3 + 6\lambda_2(\phi-1)\rho^2 + 4c\rho - 8\lambda_2(\phi-1) \end{bmatrix}\phi\beta}{\begin{bmatrix} (\phi-1)^2(\rho+1)^2\beta^4 + 2(\rho+1)(-\rho+\phi-2)\beta^2 \\ + (\rho+2)^2 \end{bmatrix}(\phi-1)\begin{bmatrix} (\phi-1)^2(\rho-1)^2\beta^4 - 2(\rho-1)(\rho+\phi-2)\beta^2 \\ + (\rho-2)^2 \end{bmatrix}}$$

2.A.2 Example: ownership matrices for three firms (each with Two Products)

Suppose there are three firms (f, f', f''), each owning two products.

- Firm f owns products 1 and 2
- Firm f' owns products 3 and 4
- Firm f'' owns products 5 and 6

The product list is $[1, 2, 3, 4, 5, 6]$. Below are the key matrices used in the matrix-form first-order conditions:

1. Full ownership Matrix $\mathcal{H}_t^{all}$

$$
\begin{array}{c c}
 & \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array} \\
\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{array} &
\begin{bmatrix}
1 & 1 & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 1 & 0 & 0 \\
0 & 0 & 1 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 \\
0 & 0 & 0 & 0 & 1 & 1
\end{bmatrix}
\end{array}
$$

2. Firm-specific matrix for f: $\mathbf{H}_t^{(f)}$

$$
\begin{array}{c c}
 & \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array} \\
\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{array} &
\begin{bmatrix}
1 & 1 & 0 & 0 & 0 & 0 \\
1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
\end{array}
$$

3. 'Other-Firm' matrix for f: $\mathbf{H}_t^{(-f)}$

$$
\begin{array}{c c}
 & \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array} \\
\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{array} &
\begin{bmatrix}
0 & 0 & 1 & 1 & 1 & 1 \\
0 & 0 & 1 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix}
\end{array}
$$

4. Matrix for products not owned by f: $\mathbf{H}_t^{-f}$

$$
\begin{array}{c c}
 & \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array} \\
\begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{array} &
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 1 & 1
\end{bmatrix}
\end{array}
$$

- $\mathbf{H}_t^{all}$ encodes which products are owned by the same firm (for all products).
- $\mathbf{H}_t^{(f)}$ isolates f's internal cross-effects.
- $\mathbf{H}_t^{(-f)}$ encodes the links between f's products and those not owned by f (used for the platform's revenue share).
- $\mathbf{H}_t^{-f}$ gives the ownership structure for products not owned by f (interactions among f' and f'' products).

2.A.3 Robustness of IV checks

Table 2.8: First-Stage IV regression results for Selected IV Set

Variables	Price		Average rating	
	Param.	SE	Param.	SE
Suction Power (KPa)	1.079***	(0.036)	−0.065***	(0.007)
Battery Life (hours)	0.911***	(0.128)	0.044	(0.027)
Height (inches)	0.011	(0.008)	−0.000	(0.001)
IV4	0.024	(0.033)	0.017*	(0.010)
IV8	0.005**	(0.002)	0.000	(0.001)
IV11	−0.014***	(0.003)	−0.002***	(0.001)
IV14	0.014	(0.017)	−0.011**	(0.005)
IV17	0.019	(0.014)	0.010***	(0.003)
Constant	−1.450	(0.984)	4.407***	(0.358)
Market FEs	Yes		Yes	
Firm FEs	Yes		Yes	
Observations	3,592		3,592	
F-stat instr (IVs)	35.89		12.24	
F p-val	0.000		0.000	

Notes: The table presents the results of first-stage IV OLS regressions for the selected instrument set. The dependent variables are price and average rating. IV4 is the number of rival products in the same market (excluding products of the focal firm); IV8 is the sum of rivals' product heights; IV11 is the sum of rivals' product weights; IV14 is the sum of rivals' battery life; and IV17 is the sum of rivals' suction power. All regressions include the same controls and market and firm fixed effects (not reported). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

2.A.4 Alternative fixed-effect specifications

Table 2.9: Robustness of joint non-linear GMM RCL estimates to alternative fixed effects

	Month × Firm FE		Month FE		Product FE		Month × Product FE		Firm FE		Quarter FE	
Variable	**Est.**	**SE**	**Est.**	**SE**	**Est.**	**SE**	**Est.**	**SE**	**Est.**	**SE**	**Est.**	**SE**
Demand Equation												
Constant	-25.243	2.898	-31.918	0.921	-152.943	6.305	-18.042	0.002	-117.228	57.970	-8.725	1.082
Suction Power	2.290	0.255	1.863	0.175	21.574	1.331	1.890	0.005	3.770	17.253	1.873	0.460
Battery Life	1.113	0.846	-0.841	0.424	-90.824	1.689	1.308	0.004	-0.012	5.261	-0.362	0.475
Height	0.004	0.038	0.026	0.035	0.837	0.252	0.010	0.001	0.017	0.291	0.062	0.028
Price ($-\alpha$)	-1.951	0.020	-2.020	0.000	-5.275	0.192	-2.482	0.000	-2.819	1.359	-1.753	0.339
Rating (τ)	1.998	0.020	5.164	0.001	59.620	1.477	0.014	0.000	25.172	7.915	0.003	0.048
Price (σ_α)	0.472	0.011	1.322	0.000	14.028	0.591	11.887	0.000	1.294	54.763	1.389	0.716
Rating (σ_τ)	0.333	0.007	2.075	0.001	28.556	1.018	2.614	0.000	0.035	42.851	3.568	0.420
Price (π_α)	-2.321	0.058	1.741	0.000	-5.148	1.703	-0.228	0.000	2.633	16.266	1.610	5.204
Rating (π_τ)	-1.060	0.045	0.924	0.000	1.487	1.286	-0.206	0.000	-2.539	9.192	-4.255	2.357
$\mathbf{c_0}$ Equation												
Constant	-24.164	3.303	-35.727	2.660	-55.249	2.798	-10.976	0.007	-67.041	49.029	-2.178	82.060
Height	0.007	0.090	0.527	0.450	1.414	0.288	0.212	0.000	0.186	1.161	0.017	1.032
Weight	0.330	0.227	0.919	0.414	0.487	0.268	0.036	0.001	-0.706	7.612	0.202	8.198
$\mathbf{c_1}$ Equation												
Constant	8.725	6.842	50.743	2.770	18.390	1.925	1.006	0.003	24.609	16.345	0.834	30.855
Suction Power	-0.199	0.192	-1.863	0.686	-2.216	0.306	-0.123	0.000	-0.643	7.863	-0.521	4.422
Battery Life	-0.484	2.072	-3.500	1.298	-1.513	0.880	-0.101	0.001	-5.879	17.791	-0.132	12.154
Model Statistics												
N	3592	-	3592	-	3592	-	3592	-	3592	-	3592	-
J-stat	2	-	32	-	0	-	54	-	6	-	60	-
N MC (<0)	20	-	22	-	2	-	54	-	2	-	25	-
Average PCM	0.261	-	0.362	-	0.119	-	0.533	-	0.184	-	0.275	-
Average $\mathbf{c_t}$	2.42	-	2.05	-	2.702	-	1.497	-	2.519	-	2.383	-
Average $\mathbf{c_t^g}$	0.0004	-	0.0014	-	0.0046	-	0.0001	-	0.0035	-	0.0001	-

Notes: The table reports GMM RCL estimates for alternative fixed-effect specifications. Each column corresponds to a different set of fixed effects included in the demand and cost equations, holding instruments and other controls fixed. Demographic variables (including income) are scaled and demeaned at the market level. 'Month × Firm FE' corresponds to the baseline specification used in the main results in Table 2.6. Standard errors are reported in separate columns.

Chapter 3

Non-price effects of M&A in the app market[1]

This chapter studies the non-price effects of mergers and acquisitions (M&A) in the US Google Play Store (2015–2019). Using 1,387 M&A involving US developers in 2016-2017 and combining matching with a staggered DiD design, we show that consolidation reallocates developer effort from portfolio expansion toward within-portfolio optimization. After M&A, developers reduce entry and external growth while increasing maintenance and incremental improvement of existing products, leading to lower portfolio turnover. These effects are driven mainly by acquisitions rather than mergers and are stronger for non-horizontal transactions, consistent with a shift toward intensive-margin incentives. Price responses are economically small and limited to few paid apps, indicating that competitive effects operate predominantly through non-price channels.

[1]This chapter is based on joint work with Michael Kummer and Franco Mariuzzo. Michael Kummer provided the raw data. I constructed the analysis dataset, defined the M&A sample, and conducted the empirical analysis using the full dataset. The empirical design and interpretation benefited from discussions and guidance from Michael Kummer and Franco Mariuzzo, who also provided comments on earlier drafts. All remaining errors are our own. This manuscript benefited from the use of AI-assisted writing tools for grammar and style improvements; all intellectual content, analysis, and interpretations remain the sole responsibility of the authors. We thank participants at the internal seminar at the School of Economics (UEA), the 3rd Workshop on Digital Economics 2024 (CMA Offices, London), *Markets and Waves 1.0* (Nova SBE, Lisbon), CLEEN 2024 (Tilburg), *Economics of Digitization 2024* (CESifo, Munich), the 18th International Conference on Competition and Regulation (MINOA Palace Resort, Chania, Crete), and the RES Annual Conference at Birmingham (30 June 2025) for helpful comments and discussions.

3.1 Introduction

Competition authorities have recently increased their focus on the impact of M&A in high-tech digital markets, particularly on innovation and other non-price factors such as product variety, quality, sustainability, and labour market outcomes (European Commission, 2019).[2] While tools for predicting the price effects of M&A are well-established, estimating non-price effects is more complex—particularly in dynamic digital markets, where their multi-sided nature adds layers of complexity. However, the rapid changes inherent to digital markets can also facilitate the identification of such effects over shorter periods (OECD, 2018).

This complexity arises from the ambiguous relationship between M&A and non-price outcomes—such as innovation and product variety—compounded by the unique traits of digital markets. For instance, young developers may act as either competitors or acquisition targets, influencing market dynamics in ways that are challenging to evaluate (Argentesi et al., 2021). Digital markets affect not only non-price dimensions of competition but also raise broader societal concerns related to cultural diversity, political pluralism, and privacy (Calvano and Polo, 2021).

This paper examines how M&A affect non-price business decisions of app developers, focusing on product variety, innovation-related activity, portfolio reallocation, and monetization choices. Using a unique dataset of 8 million apps and 1.8 million developers on the Google Play Store (2015–2019), we analyse 1,387 M&A events (126 mergers and 1,261 acquisitions) that occurred between 2016 and 2017. Our primary analysis is conducted at the developer level, with additional app-level analysis to study changes in update activity, portfolio status, and monetization outcomes.

To identify causal effects, we combine matching with a modern staggered DiD design. Specifically, we first construct a matched control group of non-M&A developers and their apps based on pre-treatment (pre-merger) characteristics, and then estimate dynamic treatment effects using the timing-of-events framework of Callaway and Sant'Anna (2021). This approach accommodates heterogeneous treatment timing, avoids the biases of two-way fixed effects estimators, and ensures

[2] For an overview of recent regulatory actions addressing price and non-price effects in digital markets, including cases involving app developers and platform practices, see subsection 3.A.1.

that post-merger outcomes are compared to credible counterfactuals with parallel pre-trends.

We contribute to the literature and policy debate by addressing key gaps in understanding the non-price effects of M&A in a digital market. Using a large dataset of 1,387 M&A within a single market, we conduct separate analyses for mergers versus acquisitions and horizontal versus non-horizontal M&A, providing detailed insights into their distinct impacts. Additionally, our developer- and app-level analyses offer a comprehensive perspective on how M&A influence digital ecosystems

Our findings show that M&A significantly reshape app-development activity by reallocating developer effort from expansion to consolidation. Following M&A, developers perform approximately 1.15 additional updates, corresponding to about 3% of a standard deviation in the population distribution of updates. By contrast, M&A reduce expansion activity: developers create around 1.72 fewer new apps and acquire 1.55 fewer purchased apps, representing about 16% and 21% of a standard deviation, respectively. M&A also reduce app variety, with decreases of 0.79 removed apps and 0.50 dismissed apps, equivalent to roughly 8% and 6% of a standard deviation in their respective distributions. Taken together, these patterns indicate that M&A shift developers away from broad portfolio expansion and toward intensified maintenance and retention of existing apps. This shift maps naturally onto the distinction between 'competition *for* the market' and 'competition *in* the market'. Before M&A, independent developers compete largely *for* markets: they explore and enter many different product niches in search of a success, and growth comes through launching and acquiring new apps. Once a developer is acquired or merges, it no longer needs to sustain this costly exploration; instead it competes *in* the markets where its successful products already operate, concentrating effort on maintaining and improving those products rather than on entering new ones. The observed combination of fewer new and purchased apps alongside more frequent updates is exactly what this reallocation from exploration toward in-market optimization would predict.

When comparing the effects of M&A, clear distinctions emerge in app development outcomes. Mergers generally exhibit positive effects on app updates and have no significant impact on app variety. This suggests a strategic emphasis on refining and consolidating existing app portfolios. In contrast, acquisitions show negative impacts on new app launches, app variety, and several app-related metrics, indicating a strategy focused on streamlining assets and operational efficiency. These contrasting effects underline the differing objectives pursued through M&A. When distinguishing between horizontal and non-horizontal M&A, the impacts diverge further. Horizontal M&A reduce app removals, dismissals, and sales, stabilizing portfolios and optimizing existing products. In contrast, non-horizontal M&A lead to fewer new app launches and purchased apps, signalling a retrenchment in variety expansion and external growth, with post-merger adjustment occurring primarily through within-portfolio development and integration. These patterns highlight that horizontal M&A prioritize portfolio stability and operational efficiency, while non-horizontal M&A tend to limit growth opportunities, reflecting their distinct strategic goals.

The structure of the paper is as follows: Section 3.2 reviews the relevant literature. Section 3.3 presents the data. Section 3.4 details the empirical strategy. Section 3.5 discusses the results at the app-developer level, with a final focus on the app-level analysis. Lastly, section 3.6 concludes with the main implications of the findings and provides a policy discussion.

3.2 Literature review

In the past decade, and particularly in more recent times, there has been a growing interest in the interplay between competition, shifts in market structure, and non-price effects such as innovation, variety, quality, employment, and sustainability. The Federal Trade Commission's 2010 Horizontal Merger Guidelines highlight the importance of innovation as a factor in merger assessments, with approximately one-third of contested mergers raising concerns related to innovation (Gilbert and Greene (2014)). The European Commission has increasingly emphasized innovation in high-profile cases such as Dow/DuPont,[3] where the merger's potential impact on agricultural R&D was scrutinized, GSK Oncology/Novartis,[4] which raised concerns about reduced innovation incentives in oncology treatments, and General

[3] `https://tinyurl.com/4easjzkc`
[4] `https://tinyurl.com/3drcv8t6`

Electric/Alstom,[5] where the focus was on preserving innovation in energy technology markets. Similarly, the UK's CMA has incorporated non-price considerations into its updated merger guidelines, reflecting growing recognition of these factors. Beyond innovation, non-price effects such as product variety, quality, and employment have also featured prominently in decisions, such as in the Amazon/Deliveroo case,[6] where labor market impacts were considered, and the Veolia/Suez merger,[7] where sustainability and environmental outcomes were critical factors. These cases demonstrate the expanding scope of competition policy in addressing non-price dimensions of mergers, particularly in dynamic and innovation-intensive sectors.

The literature on non-price effects of M&A has long examined their impact on innovation, product variety, consumer choice, and market diversity beyond price. In traditional markets, research has focused on how M&A influence innovation cycles, product quality, and employment dynamics, providing insights into industries like pharmaceuticals, manufacturing, and energy. More recently, the digitization of the economy and the rise of acquisitions involving small startups by large digital platforms have reignited interest in the unique non-price dynamics of M&A in digital markets. These studies explore how platform acquisitions affect innovation incentives, consumer choice, market diversity, and the competitive structure of multi-sided markets.

In this section, we review the literature on the non-price impacts of M&A, addressing traditional and digital markets separately. We then highlight how our study contributes to and extends this body of research.

3.2.1 M&A in traditional markets: Non-price impacts

Since the seminal works of Schumpeter J A (1942) and Arrow (1972), which explored the intersection of competition and innovation, a substantial body of literature has flourished. Noteworthy among recent contributions are studies by Federico et al. (2017), Gilbert (2019), and Bourreau et al. (2021), which theoretically examine the dynamics of market structure changes through mergers and their broader impacts on innovation and market diversity. Contemporary literature emphasizes that the effects of mergers extend beyond merely reducing the number of competitors or altering

[5] https://tinyurl.com/5b5xappb
[6] https://tinyurl.com/2spverev
[7] https://tinyurl.com/mvtmyz9p

product differentiation. Moraga-González et al. (2022) argues that mergers compel non-merging firms to adjust their innovation strategies, highlighting strategic interdependencies.

The relationship between M&A and firms' incentives to invest in R&D has been widely studied in the empirical literature. However, the effects of mergers on R&D and innovation are diverse, with studies presenting conflicting results. Szücs (2014) found that target firms tend to reduce their R&D efforts after mergers, while acquirers face lower R&D intensity due to higher sales volumes. Other empirical research, however, shows mixed outcomes (Cassiman et al. (2005), Haucap and Stiebale (2023)). For instance, Cassiman et al. (2005) analysed 31 merger cases, finding that technological and market-relatedness between M&A partners influences R&D inputs, outputs, performance, and organizational structures. In contrast, Ornaghi (2009) showed that greater technological relatedness does not necessarily lead to better R&D outcomes, as merged firms may lack incentives to boost R&D compared to non-merged firms. Similarly, Haucap et al. (2019) used PSM to demonstrate that post-merger periods often see declines in patents and R&D for both merged entities and their competitors. Bennato et al. (2021) examined mergers in the hard-drive disc industry and found that the method used to measure innovation—such as R&D intensity, new patents, or unit prices—affects results. Some studies report that mergers can lead to increased R&D activity, suggesting that the impact on innovation is context-dependent.

Mergers significantly impact product variety and consumer choice. Berry and Waldfogel (2001) found that mergers in the broadcasting industry can reduce product variety in large markets but may increase it in smaller markets by supporting offerings that were previously unviable. Sweeting (2010) showed that radio station mergers often lead to changes in product positioning, affecting consumer choice and diversity. Fan and Yang (2020) analysed the US smartphone market, highlighting how reduced competition, whether due to mergers or other factors, exacerbates consumer welfare losses by limiting variety and choice. Chipty (2001) and Crawford and Yurukoglu (2012) examined the cable television industry, noting that vertical integration can enhance efficiency and product offerings, such as channel availability, but also poses risks like market foreclosure and reduced competition. These findings emphasize that strategic realignment of products is a significant non-price effect of mergers, with nuanced outcomes depending on market structure and consumer dynamics. Johnson and Rhodes (2021) underscored the importance of product quality in merger analysis, introducing the product-mix effect as a crucial factor influencing consumer surplus. This challenges conventional merger assessment tools like the HHI and calls for a more nuanced approach when evaluating mergers in markets where product quality plays a significant role.

Labour market outcomes represent a critical but often under-explored non-price impact of mergers. Marinescu and Hovenkamp (2019) investigated how mergers affect labour market monopsony power, demonstrating that increased employer concentration can suppress wage growth and limit job mobility. Prager and Schmitt (2021) provided empirical evidence from the hospital sector, showing that employer consolidation slows wage growth for skilled workers. Neven and Röller (1996) examined the European airline industry, finding that reduced competition leads to significant rent sharing through higher wages. These studies highlight the broader economic and societal implications of mergers, extending the analysis beyond traditional price and output effects.

Environmental sustainability has also become a crucial concern in merger evaluations. OECD (2021) identified four primary theories of harm: increased market power leading to higher prices and reduced green quality, buyer power dampening incentives for green innovation, unilateral effects slowing innovation, and 'green killer acquisitions' targeting sustainable competitors to reduce competitive pressure. Fikru and Gautier (2020) analysed high-polluting industries, highlighting the moderating role of external factors like emission taxes in shaping the environmental impact of mergers. Similarly, Liang et al. (2022) found that mergers in Chinese heavy-polluting industries could foster green innovation when supported by government subsidies. These studies underscore the complex interplay between mergers and sustainability, advocating for nuanced regulatory frameworks.

While extensive, the current literature on non-price effects in traditional markets often overlooks the granularity of how mergers impact firm-level strategic decisions and externalities. Most studies focus on traditional markets with physical products, neglecting digital markets, where products can be easily modified or new ones quickly introduced. Our study aims to fill this gap by providing empirical evidence on the differentiated non-price outcomes of horizontal versus non-horizontal mergers in digital markets, particularly regarding product quality, innovation, and product variety.

3.2.2 M&A in digital markets: Non-price impacts

The digitization of the economy has added new dimensions to M&A impacts, particularly concerning non-price factors such as product quality, consumer choice, data privacy, and market diversity. Major tech companies like GAFAM have been active in acquisitions, often escaping antitrust scrutiny. This has raised concerns about market concentration and power in digital markets, prompting high-profile reports European Commission (2019), which call for better competition promotion (see Motta and Peitz (2021)).

Researchers such as Letina et al. (2024) emphasize that a permissive approach to merger control often overlooks the elimination of potential competition, particularly in digital markets where innovation and technological synergies are key drivers of competitive dynamics. Etro (2019) analyses a merger between two digital platforms, focusing on network effects and the potential loss of platform variety, using the Rover and DogVacay case study.[8] Gautier and Lamesch (2021) provides a comprehensive analysis of GAFAM's acquisition strategies, identifying key characteristics of the firms they target. The structural differences between platforms and traditional vertically integrated firms, along with the significance of intangible assets like data, are highlighted by Parker et al. (2021).

Despite extensive research on the effects of M&A on innovation,[9] there remains limited evidence on how M&A affect non-price developer behaviour and product-portfolio dynamics within app-based digital ecosystems. Existing literature suggests that GAFAM acquisitions often reduce innovation among competing apps. Affeldt and Kesler (2021) finds that such acquisitions diminish innovation efforts of competing apps without significantly affecting app prices or privacy permissions. Additionally, Kesler et al. (2019) shows that increased market concentration in digital app platforms correlates with heightened data collection, raising privacy concerns. This suggests that greater market power may lead to reduced consumer privacy, a crucial non-price effect of M&A.

[8] `https://tinyurl.com/mrc275av`

[9] We also analyse separately the acquiring behaviour of GAFAM. Refer to subsection 3.A.2 for details.

In the mobile app ecosystem, Ghose and Han (2014) reveals that in-app purchases boost demand, while in-app advertisements tend to reduce consumer interest. These monetization strategies shape competitive dynamics in app markets. Yin et al. (2014) explores innovation paths in the iPhone ecosystem, highlighting how different app categories, such as games versus non-games, follow distinct innovation trajectories. This reflects the diversity of innovation within app markets. Moreover, Janßen et al. (2022) shows that privacy regulations, such as General Data Protection Regulation (GDPR), lead to market exits and decreased new app entries, indicating a trade-off between regulatory compliance and market diversity.

Collectively, these studies emphasize the significant non-price effects of M&A in digital markets. They demonstrate how acquisitions influence innovation, market concentration, consumer privacy, and product diversity. These non-price factors shape competitive dynamics in the app ecosystem, offering a deeper understanding of the evolving digital economy beyond traditional price-based analyses.

3.2.3 Contribution to the literature

This study makes several key contributions to the literature. First, it offers an in-depth analysis of the non-price effects of M&A involving app developers in digital markets, focusing on product variety, launch rates, sales dynamics, and app updates. Unlike previous studies such as Szücs (2014), Haucap et al. (2019), and Stiebale and Szücs (2022), our research distinguishes between M&A, recognizing that they have different effects on market behaviour.

Second, our empirical analysis reveals heterogeneous effects between horizontal and non-horizontal M&A. Horizontal M&A are associated with reductions in app removals and sales, indicating greater portfolio stability following consolidation. By contrast, non-horizontal M&A more strongly curb expansionary margins, including new app launches and external app purchases. This distinction sharpens our understanding of the different strategic objectives underlying these transactions and highlights how their non-price effects operate through portfolio composition, development intensity, and product variety rather than prices.

Third, our findings highlight the distinct strategies adopted by merging firms in response to market conditions. Horizontal M&A prioritize maintaining the installed base, while non-horizontal M&A emphasize improving product quality through more frequent app updates. This adds to the literature by demonstrating the importance of differentiating M&A types when assessing their impact on consumer choice and market diversity.

Additionally, we contribute to the literature on digital market dynamics, where the relationship between product variety, updates, and market concentration differs from that in traditional industries. Using a unique dataset of 1,387 M&A transactions in the US Google Play market, we provide empirical evidence that non-price effects—such as changes in product portfolios and strategic app updates—are crucial for understanding the broader implications of M&A on consumer welfare beyond pricing. To assess the external validity of these effects, we additionally replicate the analysis on Google Play data for the Japanese market as a robustness check (Table 3.20), finding consistent patterns across the two markets.

Finally, we contribute to the literature on product repositioning in digital markets. Building on studies of traditional sectors such as radio and airlines (Berry and Waldfogel (2001); Li et al. (2022)), we examine how M&A affect product variety and within-platform portfolio adjustment in digital environments, where the flexibility of digital products allows for rapid change. Using data from the US Google Play Store, we provide evidence on how M&A reshape product offerings and key non-price dimensions of competition in app markets.

3.3 Data

Mobile applications have become indispensable in modern life, primarily distributed through two dominant platforms: the Apple App Store and Google Play Store. In 2020, global mobile app spending surged to $327.1 billion-an impressive 17% growth over the previous year-and drew a record $120 billion in investments.[10] Google Play alone boasts a massive user base of 2.5 billion active users, highlighting the immense opportunities for developers in this thriving ecosystem.[11] Developers leverage personalized data to create user-focused apps or monetize through advertising, fuelling an app development market valued at $197.2 billion in 2021.[12] This dynamic

[10] https://tinyurl.com/3knv33xd

[11] https://blog.gitnux.com/google-play-store-statistics/

[12] https://straitsresearch.com/report/mobile-app-development-market/

market serves as the backdrop for our study, which examines the strategic behaviours of app developers engaged in M&A.

To support our analysis, we compiled a comprehensive dataset from the US Google Play Store, covering both product-level and developer-level information. Collected at six-week intervals from 2015 to 2019, this dataset enables us to track M&A activities by identifying changes in developer ownership structures.[13]

3.3.1 Summary statistics - developer Level

Before turning to the regression analysis, we define the two groups that we compare throughout the chapter. We refer to *M&A developers* as the 1,387 US developers that complete at least one merger or acquisition during the 2016–2017 event window—this is our treated group, taking each developer's first M&A event. We refer to *non-M&A developers* as the remaining developers in the panel that never engage in M&A over the 2015–2019 sample period, which form the comparison pool from which controls are later matched. Table 3.1 compares the two groups across outcome and control variables, pooling all developer-period observations. As the table makes clear, the two groups differ substantially at baseline, which is precisely why the causal analysis below relies on matching rather than on a raw comparison of means.

Developers involved in M&A are systematically more active and diversified than their non-M&A counterparts. As shown in Table 3.1, the average M&A developer releases 1.20 new apps per period, compared to only 0.18 among non-M&A developers—a more than sixfold difference. External expansion is also far more pronounced: M&A developers purchase an average of 0.41 apps from others, versus 0.04 for non-M&A firms, nearly a tenfold increase. Even when focusing on internally generated 'newborn' apps, the gap remains substantial, with M&A developers producing 0.79 compared to 0.29. These patterns suggest that M&A developers pursue a dual strategy of internal innovation and external acquisition.

Quality-enhancing activities also differ markedly. M&A developers conduct 2.23 updates on average, compared to 0.29 among non-M&A developers, nearly eight times as many. The distinction is particularly sharp for Minor updates (1.67 versus 0.24), while Major updates, though less frequent, still occur more than ten times as

[13]Details on defining and identifying M&A are provided in subsection 3.A.3.

often among M&A developers (0.56 versus 0.04). This pattern indicates that M&A developers devote systematically greater resources to refining and maintaining their products, consistent with a strategy aimed at user retention and quality improvement.

Portfolio optimization further differentiates the two groups. M&A developers remove 0.71 apps compared to 0.12 for non-M&A developers, dismiss 0.55 versus 0.11, and sell 0.19 versus 0.09. Although the absolute numbers are small, the proportional differences are striking: M&A developers are five times more likely to dismiss apps and more than twice as likely to sell them. Such selective pruning and reallocation reflects active portfolio management and suggests that M&A facilitate redeployment of resources toward higher-value projects.

Monetization strategies also diverge. M&A developers adjust pricing schemes more frequently, with a mean of 0.09 monetization changes compared to 0.02 for non-M&A developers. Yet they sustain a smaller share of paid apps (1% versus 5%), indicating a tilt toward freemium or indirect monetization models. Together, these patterns suggest that M&A developers prioritize scale and engagement over upfront revenues, leveraging larger user bases to monetize through in-app purchases or advertising.

Control variables confirm that these differences are not incidental. M&A developers own, on average, more than four times as many apps (16.47 versus 3.86) and have accumulated far more ratings (about 151,686 versus 4,194 on average; the very large means and standard deviations reflect a small number of high-profile developers with heavily skewed rating counts, which is why these variables enter the matching and regressions in logs). They are also more internationally oriented, holding more apps with translated descriptions (3.55 versus 1.11). They display slightly higher baseline quality, with average ratings of 3.60 compared to 3.47. These figures highlight the pre-existing scale and sophistication of M&A developers, underscoring that selection into treatment is non-random.

Taken together, the descriptive statistics reveal systematic differences between M&A and non-M&A developers along dimensions of innovation, portfolio management, and monetization. These patterns document the proactive strategies pursued by M&A developers but also highlight the need to account for these difference in our empirical strategy. Because treated firms are larger and more dynamic prior to M&A, we will employ PSM combined with a DiD framework to construct credible counterfactuals when isolating the causal impact of consolidation on app market outcomes.

3.3.2 Summary statistics - app level

At the app level, systematic differences emerge between products associated with M&A developers and those of non-M&A developers. As shown in Table 3.2, apps in the M&A group are updated more frequently: the probability of an update is 0.16 compared to 0.07 for non-M&A apps, more than doubling the baseline rate. This gap underscores the greater emphasis that M&A developers place on sustaining and improving product quality through ongoing refinements.

Differences in user feedback and visibility further reinforce this pattern. M&A apps attain an average rating of 3.58, compared to 3.11 for non-M&A apps, a non-trivial difference given the five-point scale. Moreover, the volume of user feedback is substantially higher: the log of ratings averages 3.12 for M&A apps versus 2.17 for non-M&A apps, implying that M&A apps receive more than twice as many ratings in raw terms. These disparities suggest that M&A apps not only attract more users but also elicit stronger engagement, reflecting both greater market penetration and sustained consumer attention.

Portfolio dynamics also differ in important ways. M&A apps are slightly less likely to be dismissed in the subsequent period (3.6% versus 5.7%) but are three times as likely to be sold to another developer (1.8% versus 0.6%). The elevated sale rate indicates that M&A developers are more active in reallocating app assets across publishers, consistent with a market in which ownership restructuring plays a key role in shaping outcomes. By contrast, the higher dismissal rate among non-M&A apps suggests weaker resilience, with a larger share of their products failing to remain viable in subsequent periods.

Price distributions diverge as well. The mean listed price for M&A apps is \$0.20, markedly lower than the \$0.34 observed for non-M&A apps, even though the variance in prices is much greater for the latter. This contrast highlights a strategic tilt among M&A developers toward low-cost or freemium pricing models designed to maximize user acquisition, in contrast with the broader and more dispersed pricing practices of non-M&A developers.

Table 3.1: M&A vs Non-M&A summary statistics 2015-19

	(M&A)		(Non M&A)		(Mean diff)	
	Mean	SD	Mean	SD	Coeff.	(t-value)
Outcome variables						
(1)#New apps	1.20	10.58	0.18	2.50	1.02[a]	(21.34)
(1a)#Purchased apps	0.41	7.41	0.04	2.32	0.38[a]	(11.25)
(1b)#Newborn apps	0.79	7.48	0.29	2.87	0.50[a]	(14.63)
(2)#App updates	2.23	34.39	0.29	5.37	1.94[a]	(12.46)
(2a)#Major updates	0.56	25.40	0.04	1.22	0.51[a]	(4.46)
(2b)#Minor updates	1.67	23.08	0.24	5.18	1.43[a]	(13.66)
(3)#Removed apps	0.71	9.37	0.12	2.49	0.59[a]	(13.92)
(3a)#Dismiss apps	0.55	7.71	0.11	2.43	0.44[a]	(12.45)
(3b)#Sold apps	0.19	0.56	0.09	0.34	0.09[a]	(37.14)
(4)Share paid apps	0.01	0.08	0.05	0.19	-0.04[a]	(-108.89)
(5)ΔMonetization strategy	0.09	1.21	0.02	0.97	0.07[a]	(12.52)
Control variables						
#Owned apps	16.47	84.25	3.86	20.82	12.61[a]	(33.01)
#Ratings (thousands)	151.69	3,051.19	4.19	229.04	147.49[a]	(10.66)
#Paid apps	0.80	5.46	0.29	5.35	0.50[a]	(20.27)
#Translated des	3.55	12.76	1.11	8.47	2.44[a]	(42.13)
#Miss downloads	0.22	6.31	0.06	2.07	0.16[a]	(5.76)
Average rating	3.60	0.99	3.47	1.56	0.13[a]	(28.95)
#Status updates	3.22	36.90	0.56	6.51	2.66[a]	(15.92)
N observations	48,597		23,471,983		23,520,580	

Notes: The superscript [a] denotes statistical significance at the 1% level. This table reports summary statistics for developers involved in mergers or acquisitions (M&A) and those who were not (Non-M&A). All variables are reported in levels. The control variables (*#Owned apps*, *#Ratings*, *#Paid apps*, *#Translated des*, and *#Miss downloads*) enter the propensity-score matching and the regressions in logs (using $\log(1 + x)$ for the count variables) to limit the influence of their heavy right tails, but are shown here in levels for transparency. The variables are defined as follows: *#Owned apps* is the number of apps currently owned by the developer, and *#App updates* counts the total number of updates made to existing apps during the observation period. *#Major updates* refers to updates that introduce substantial feature additions or redesigns, as identified from update-note text, while *#Minor updates* captures small-scale maintenance changes such as bug fixes or stability improvements. *#Removed apps* measures the number of apps withdrawn from the platform, while *#New apps* counts newly introduced apps. *#Ratings* is the cumulative number of ratings received by the developer's apps, reported in thousands, and *Average rating* is the mean user rating across these apps. *#Sold apps* denotes the number of apps transferred to another developer, whereas *#Dismiss apps* captures apps entirely removed from the store. *#Purchased apps* measures the number of apps acquired by the developer. *Share paid apps* is the proportion of apps that are paid rather than free. Δ*Monetiz. strategy* records the number of apps whose monetization strategy changed (e.g., price adjustments or switches between free and paid). *#Translated des* is the number of apps with translated descriptions, and *#Miss downloads* is the number of apps lacking download information. Finally, *#Status updates* records changes in an app's availability or standing within Google Play (e.g., published, suspended, removed, restored), as tracked by AppMonsta.

Table 3.2: Summary statistics - app level

Statistic	N	Mean	St. Dev.	Min	Max
M&A					
App update	453,101	0.16	0.36	0	1
Price($)	453,101	0.20	1.98	0	135.75
Dismissed following period	446,173	0.036	0.187	0	1
Sold following period	429,970	0.018	0.131	0	1
Miss downloads	453,101	0.005	0.07	0	1
Average rating	453,101	3.58	1.61	0	5.00
Log(1+#Ratings)	453,101	3.12	2.70	0	17.98
Status update	453,101	0.18	0.38	0	1
Non M&A					
App update	71,862,291	0.07	0.25	0	1
Price($)	71,862,291	0.34	5.45	0	746.41
Dismissed following period	69,604,321	0.057	0.231	0	1
Sold following period	65,652,594	0.006	0.075	0	1
Miss downloads	71,862,291	0.01	0.10	0	1
Average rating	71,862,291	3.11	1.93	0	5.00
Log(1+#Ratings)	71,862,291	2.17	2.17	0	18.25
Status update	71,862,291	0.09	0.29	0	1
M&A: #Apps : 23131; Non M&A: #Apps : 6209697					

Notes: This table presents summary statistics at the app level for apps that experienced a merger or acquisition (M&A) and those that did not (Non-M&A) from 2015 to 2019. *App update* is a binary indicator for whether the app was updated in the period. *Price ($)* is the listed app price in USD. *Dismissed following period* indicates whether the app has been dismissed. *Sold following period* is a binary variable indicating whether the app is sold (i.e., the publisher changes) in the next period. *Miss downloads* indicates missing download count information. *Average rating* is the app's cumulative average rating as reported by Google Play, and *Log(1+#Ratings)* is the natural logarithm of one plus the total number of ratings. The sample covers 23131 apps in the M&A group and 6209697 in the Non-M&A group. *#Status updates* records changes in an app's availability or standing within an Google play store, such as when an app is published, removed, suspended, or restored. This variable tracks the lifecycle events of an app as observed and recorded by AppMonsta.

Finally, missing download information-although rare overall-appears less frequently among M&A apps (0.5% versus 1%), suggesting either superior reporting quality or stronger alignment with platform requirements. Taken together, these differences indicate that M&A developers maintain apps that are better rated, more actively updated, and more commercially traded than those of their non-M&A peers.

From an empirical perspective, these contrasts emphasize the importance of accounting for pre-existing differences between M&A and non-M&A apps when estimating treatment effects. The higher baseline update probability, stronger user engagement, and distinct ownership dynamics observed in the M&A group highlight potential selection into treatment, making matching and DiD approaches essential for credible identification of causal effects.

3.3.3 Trends in key developer outcomes

Figure 3.1 displays time-series patterns for key app developer outcomes that form the basis of our empirical analysis. Several activity measures exhibit pronounced fluctuations over the sample period. In particular, the number of new apps and newborn apps rises steadily through 2017 and peaks in early 2018, followed by a subsequent decline. A similar pattern is observed for removed apps, which increase sharply around 2018, consistent with intensified portfolio restructuring during that period.

Update activity follows a distinct trajectory. App updates increase throughout 2017 and reach a clear peak in 2018, suggesting a broad shift toward maintenance and incremental improvement of existing products. This pattern coincides with major regulatory and institutional changes affecting app developers, most notably the introduction of the GDPR in 2018, which plausibly increased compliance-related update activity. We do not attribute these aggregate movements causally to regulation, but note that they represent an important common shock affecting developer behaviour.

By contrast, purchased apps and sold apps display relatively stable dynamics over time, with no comparable surge around 2018. Dismissed apps-defined as apps permanently removed from the platform-exhibit a noticeable increase in early 2018, further indicating heightened portfolio adjustment during this phase. Finally, the number of apps owned by developers follows an inverted U-shaped pattern, expanding up to 2017 before declining thereafter, consistent with a transition from portfolio expansion to consolidation.

Motivated by these patterns, our baseline analysis focuses on M&A events occurring between mid-2016 and late-2017. This window ensures a sufficiently long and relatively stable pre-treatment period while avoiding direct overlap with the most pronounced market-wide disruptions observed in 2018. The event-study specifications nonetheless exploit the full sample period, with time fixed effects absorbing common shocks such as regulatory changes.

3.3.4 Mergers and acquisitions

Table 3.3 provides a six-weekly count of M&A in the app development industry, with data points spanning from April 2016 to January 2018. For each event date, the table displays the total number of M&A combined, as well as separate counts for M&A. In our main analysis, we align with the majority of the literature by focusing only on the first M&A event for each developer. We include only those M&A with at least eight pre-event periods, ensuring a robust baseline and minimizing potential biases from earlier activities. The total number of M&A is 1,387.[14] This baseline approach enhances the validity of our results by providing a sufficient time window for pre-event observations and controlling for developer activity prior to the M&A event.

3.3.5 Market definition using app categories

We construct category-based market groupings and classify transactions using US Google Play categories under a strict, pre-event notion of developer scope. Let t_0 denote the calendar date of an M&A event. We write t_0^- to denote the *state of the world strictly before the event*: all information dated earlier than t_0 is admissible, while any observation on or after t_0 is excluded. Operationally, for each event dated t_0, we use the most recent data snapshot whose timestamp is strictly less than t_0. This procedure prevents look-ahead bias and ensures that transaction classification reflects pre-event developer activity only.

[14]Without imposing restrictions, the sample includes 2,177 M&A events for developers during the period (233 mergers and 1,944 acquisitions), as detailed in subsection 3.A.4. Results for samples with 4 and 6 pre-event observation periods are provided in subsection 3.A.5

Figure 3.1: Trends in key app developer metrics

Notes: The data presented the average value for new apps, updates, removed apps, and the number of apps owned by the developer level.

Table 3.3: Summary of M&A by Event Date

No.	Event Date	Sum MA	Count Merger	Count Acquisition
1	2016-02-15	81	7	74
2	2016-04-04	97	7	90
3	2016-05-16	69	6	63
4	2016-07-04	81	8	73
5	2016-08-15	66	5	61
6	2016-10-03	79	7	72
7	2016-11-14	92	6	86
8	2017-01-02	78	7	71
9	2017-02-13	77	6	71
10	2017-04-03	105	17	88
11	2017-05-15	66	6	60
12	2017-07-03	99	14	85
13	2017-08-14	89	1	88
14	2017-10-02	93	11	82
15	2017-11-13	117	11	106
16	2018-01-01	98	7	91
	Sum	1387	126	1261

Note: The event date represents the number of M&A that occurred in the preceding six week interval, recalling that, in our dataset, each period covers a span of six weeks. For example, the event date '2016-04-04' refers to all M&A that took place between '2016-02-15' and '2016-04-04'. The distinction between M&A lies in their outcomes: a merger involves the combination of two app developers to create a new entity, while an acquisition occurs when one app developer takes over another, resulting in the latter ceasing to exist as an independent entity. The earliest M&A activity within the observation window is used to define the event date. This dataset also includes M&A initiated by GAFAM companies. For a separate analysis of GAFAM, refer to subsection 3.A.2.

At the app level, we construct a monotone stock of cumulative ratings and retain, for each app, its maximum stock observed prior to t_0 (breaking ties by the latest pre-t_0 date). Aggregating these app-level stocks to the developer–category level yields a pre-event distribution of activity across categories. A developer's dominant category at t_0^- is defined as the category accounting for the largest share of its pre-event activity. Developers with no pre-event mass are left unclassified for that event.

Markets are defined on a dominant-category basis: only developers whose dominant category at t_0^- coincides with a given category are treated as participants in that category at t_0. This construction is used solely to classify overlap between transaction parties and does not rely on post-event information or assumptions about demand-side substitutability.

Within this framework, we classify an M&A as horizontal if, at t_0^-, the surviving developer and all counterparties share the same dominant category. If any counterparty differs, or if any party lacks a valid dominant category at t_0^-, the transaction is classified as non-horizontal. This rule is applied uniformly across events. Because mergers are relatively infrequent in our data, we pool mergers and acquisitions and report results by horizontal versus non-horizontal M&A. The strict reliance on t_0^- anchors classification before the event and avoids post-event relabelling that could mechanically inflate measured overlap.

Table 3.4: Horizontal M&A versus Non-Horizontal M&A: entity distribution

Horizontal M&A	502
Non-Horizontal M&A	885
Sum	1387

Notes: Counts reflect our pooled sample of M&A classified using pre-event dominant categories. An event is horizontal only if the survivor and all counterparties share the same dominant category at t_0^-; missing dominance renders the event non-horizontal by construction. This classification is conceptually separate from concentration metrics.

3.4 Empirical strategy

We employ an event study analysis within a DiD framework to evaluate the impact of M&A on various outcomes. This approach can also allow for the events being scattered over different times. It involves comparing changes in outcomes over time between a treatment group (entities affected by the event) and a control group (entities not affected by the event). By examining data from periods before and after the event, the DiD framework helps control for external factors that might influence the outcomes. Provided that we can find a suitable control group, the approach allows isolating the effect of the event itself. Identification in a DiD model requires that the treated and control units exhibit parallel trends prior to the event.

3.4.1 PSM and balance check

Our empirical strategy involves identifying suitable control units for the treated M&A entities. Given the large number and heterogeneity of potential controls, we use PSM to enhance comparability between treated and control observations.

While PSM is primarily designed for cross-sectional data, we adapt it for our panel dataset by converting the panel into cross-sectional data for each cohort of treated units. This approach allows us to manage the imbalance between the large control group and the smaller treatment group, which limits frequent matches in a time-series framework. For the analysis, we include app developers with a minimum of eight observed periods before the M&A event, ensuring sufficient pre-treatment data for matching. We focus on mergers from 2016 to 2017 to (a) ensure eight pre-treatment periods (each period spans six weeks) and (b) avoid external shocks like GDPR, as discussed earlier. Results for four- and six-period lags are included in the appendix.

Table 3.5 compares the control and treatment groups across various app-related metrics for four- and eight-period lags before M&A events (see subsection 3.A.6 for details on all lags). The balance check shows that matching reduces differences between the groups for most variables, though some outcome variables, such as new apps, removed apps, dismissed apps, and partially sold apps, still display significant mean differences. These discrepancies likely reflect the challenge of meeting all matching criteria in a large sample using PSM. Additional results for other lags are available in the appendix.

Table 3.5: Balance check after matching

PSM(1:1)	Control	Treatment	Mean-diff	P-value	N
Outcomes					
#New apps (lead 4)	0.52	0.93	0.411	**0.01**	2697
#New apps (lead 8)	0.47	1.01	0.539	**0.00**	2697
#Purchased apps (lead 4)	0.18	0.26	0.072	0.22	2697
#Purchased apps (lead 8)	0.24	0.32	0.084	0.22	2697
#Newborn apps (lead 4)	0.34	0.68	0.339	**0.03**	2697
#Newborn apps (lead 8)	0.40	0.69	0.292	0.08	2697
#App updates (lead 4)	1.28	1.62	0.340	0.32	2697
#App updates (lead 8)	1.24	3.55	2.316	0.16	2697
#Removed app (lead 4)	0.15	0.29	0.142	**0.00**	2697
#Removed app (lead 8)	0.13	0.26	0.132	**0.02**	2697
#Dismissed apps (lead 4)	0.13	0.23	0.100	**0.00**	2697
#Dismissed apps (lead 8)	0.11	0.21	0.106	**0.04**	2697
#Sold apps (lead 4)	0.02	0.06	0.042	**0.01**	2697
#Sold apps (lead 8)	0.02	0.05	0.026	0.12	2697
Controls					
Log(#Owned apps) (lead 4)	1.72	1.75	0.035	0.43	2697
Log(#Owned apps) (lead 8)	1.53	1.52	-0.007	0.87	2697
Log(1+#Ratings) (lead 4)	6.30	6.49	0.192	0.10	2697
Log(1+#Ratings) (lead 8)	5.91	6.03	0.121	0.30	2697
Average rating (lead 4)	3.63	3.68	0.051	0.15	2697
Average rating (lead 8)	3.65	3.67	0.021	0.59	2697
Log(1+#Paid apps) (lead 4)	0.16	0.20	0.036	0.10	2697
Log(1+#Paid apps) (lead 8)	0.16	0.19	0.033	0.12	2697
Log(1+#Translated des) (lead 4)	0.68	0.71	0.033	0.38	2697
Log(1+#Translated des) (lead 8)	0.62	0.63	0.008	0.81	2697
Log(1+#Miss downloads) (lead 4)	0.05	0.07	0.015	0.19	2697
Log(1+#Miss downloads) (lead 8)	0.05	0.06	0.015	0.18	2697
#Status updates (lead 4)	1.74	2.55	0.814	0.05	2697
#Status updates (lead 8)	2.25	4.51	2.253	0.22	2697

Notes: This table reports balance checks for treated and control units after 1:1 matching using Mahalanobis distance. The unit of observation is the developer–event, evaluated at event-time leads four or eight periods prior to the M&A date t_0. Treated observations correspond to the 1,387 M&A events in our sample. Control observations are selected from the pool of never-treated developers. Matching is performed *with replacement*, so the same control developer can be matched to multiple treated events. As a result, the matched sample contains 1,387 treated observations but only 1,310 *distinct* matched control observations (some controls appear multiple times in the matched pairs). The column N reports the total number of developer–event observations used in the balance test at the corresponding lead (treatment plus matched controls, counting repeated controls when they are reused). Reported values are group means, mean differences (treatment minus control), and p-values from two-sided tests of equality. Variables with p-values below 0.05 are highlighted in bold. Balance checks for additional leads are reported in subsection 3.A.6; variable definitions are provided in Table 3.1.

3.4.2 CSDID estimator

M&A, our treatment units, occur at different points in time. Recent econometric research highlights limitations of Two-Way Fixed Effects (TWFE) estimators when treatment timing varies, and treatment effects are heterogeneous. Even under valid parallel trends and no-anticipation assumptions, TWFE estimators can produce biased results due to incorrect weighting of treatment effects, including negative weights (de Chaisemartin and D'Haultfœuille (2020); Callaway and Sant'Anna (2021)). Bias arises because earlier-treated groups are used as controls, which can introduce spillover effects.

Given our dataset contains a large number of never-treated units, TWFE's limitations are further exacerbated. We adopt the methodology of Callaway and Sant'Anna (2021) (CSDID) to address these challenges:

$$Y_{i,t} = \sum_{e=-K}^{-2} \delta_e MA_{i,t}^e + \sum_{e=0}^{L} \beta_e MA_{i,t}^e + X_{i,t}\tau + \alpha_i + \gamma_t + \epsilon_{i,t} \tag{3.1}$$

Here, $MA_{i,t}^e$ is a dummy indicating whether unit i is e periods away from initial treatment at time t, where K represents the number of pre-treatment periods (eight in our baseline analysis) and L the post-treatment period. $X_{i,t}$ includes time-varying observed characteristics with coefficients τ; α_i and γ_t are group and time fixed effects, and $\epsilon_{i,t}$ is the error term. $Y_{i,t}$ measures app developer i's outcomes (e.g., new, updated, removed, or total apps at time t).

The term $\sum_{e=-K}^{-2} \delta_e MA_{i,t}^e$ captures the effects of pre-treatment periods, with δ_e as the coefficients. The standard practice is to omit $MA_{i,t}^{-1}$, using it as the reference period. Post-treatment effects are captured by $\sum_{e=0}^{L} \beta_e MA_{i,t}^e$, with β_e representing the coefficients for e-periods after treatment..

We focus on the Average Treatment Effect on the Treated (ATT) rather than the Average Treatment Effect (ATE). Our objective is to quantify the impact of M&A on developers that actually undergo consolidation, which is the policy-relevant estimation in this setting. Treatment is highly selective, and developers involved in M&A differ systematically from never-treated developers along multiple dimensions even prior to the event.

Moreover, treatment effects are likely heterogeneous across developers and over time. In such contexts, the ATE would combine effects for treated units with counterfactual effects for developers that are structurally unlikely to engage in M&A, requiring extrapolation beyond the region of common support. By contrast, focusing on the ATT allows us to estimate the causal effect of M&A for the population of interest without imposing assumptions about outcomes for never-treated developers under treatment.

Consistent with this objective, we combine propensity-score matching with the CSDID, constructing a comparable control group and estimating group-time average treatment effects for treated developers. This approach isolates post-M&A changes relative to credible counterfactual paths and avoids the weighting pathologies associated with two-way fixed effects estimators under heterogeneous treatment timing.

3.5 Results

3.5.1 Main results at the app developer-level

The results in Table 3.6, based on the specification with eight pre-treatment periods, reveal pronounced changes in developer behaviour following M&A. Relative to their counterfactual activity absent the acquisition, treated developers create on average 1.723 fewer new apps per period and acquire 1.552 fewer apps from other developers. New app creation and external app purchases are extensive-margin measures because they capture growth of the portfolio through the entry of additional products—either internally developed or acquired from other developers—rather than changes in the intensity of effort on apps the developer already owns; a decline in both therefore represents a contraction in the expansion and growth margins of the portfolio. These effects point to a marked contraction in extensive-margin expansion and a shift away from growth through new development or external acquisitions.

Instead, post-merger adjustment operates primarily through intensified management of existing products. Developers increase update activity by 1.145 updates per developer-period,[15] indicating greater attention to maintenance,

[15]To better understand the nature of increased update activity following M&A, we distinguish between key' and small' updates based on a structured classification of release note content, as detailed in Appendix 3.A.10.

refinement, and incremental quality improvements. At the same time, portfolio turnover declines: developers remove 0.789 fewer apps, dismiss 0.504 fewer apps, and sell 0.285 fewer apps per period relative to the counterfactual. Taken together, these patterns suggest that consolidation is associated with portfolio stabilization and internal restructuring rather than active reallocation through entry, exit, or ownership transfers.

Developers also adjust their monetization strategies following M&A. While the share of paid apps declines modestly by 0.005, the positive effect on changes in monetization strategy (ΔMonetiz. strategy = 0.046) indicates active experimentation along revenue models, such as switching between paid and free formats or refining pricing and in-app purchase structures. Overall, the evidence points to a post-merger strategy centred on improving and retaining existing apps, with comparatively limited emphasis on portfolio expansion or market-based reallocation.

Table 3.6: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT	**-1.723***	**-1.552***	-0.045	**1.145****	**-0.789*****	**-0.504*****	**-0.285*****	**-0.005*****	**0.046****
	(0.918)	(0.914)	(0.0888)	(0.514)	(0.196)	(0.166)	(0.0947)	(0.00196)	(0.0211)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	73876	73876	73876	73876	73876	73876	73876	73876	
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: The table reports the ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8 pre-treatment periods, where the matching is applied to units with at least 8 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The event study results (Figure 3.2 and Figure 3.3) corroborate this strategic shift. Following M&A, developers reduce the launch and acquisition of new apps (Figure 3.2a, Figure 3.2b), focusing instead on app updates (Figure 3.2d) and retention. The decline in removed, dismissed, and sold apps (Figure 3.3a, Figure 3.3b, Figure 3.3c) further supports this post-merger emphasis on internal consolidation.

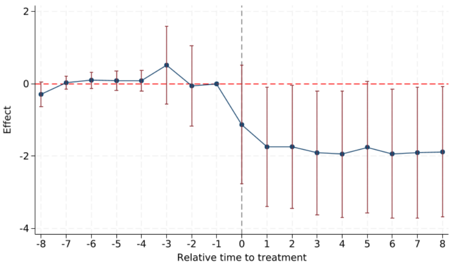

(a) #New apps

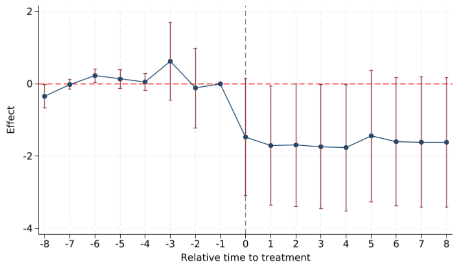

(b) #Purchased apps

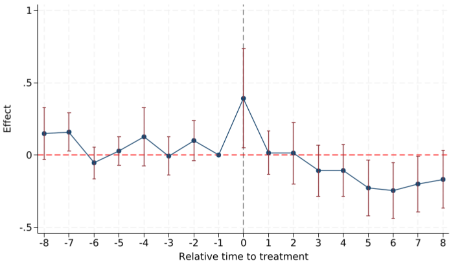

(c) #Newborn apps

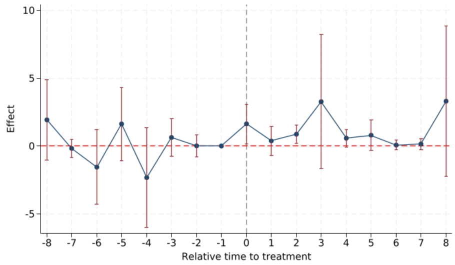

(d) #App updates

Figure 3.2: Comparative event study: Impact of M&A on app development (Part 1)

Notes: This figure presents the event study results showing the effect of M&A on different app-related outcomes. The x-axis represents the number of dates relative to the M&A event, with 0 marking the six week interval after the acquisition.

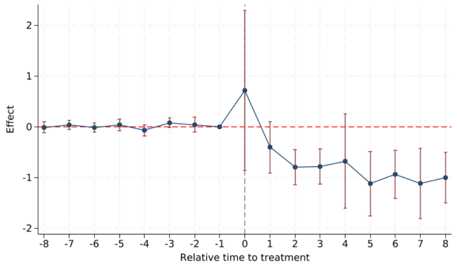

(a) #Removed apps

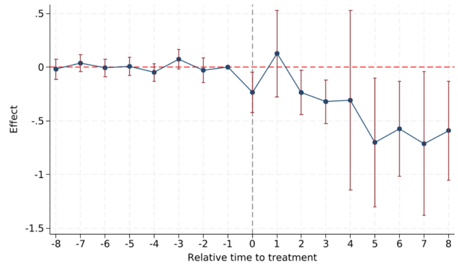

(b) #Dismiss apps

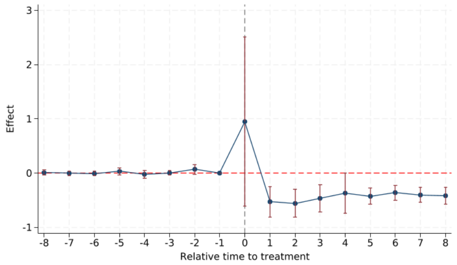

(c) #Sold apps

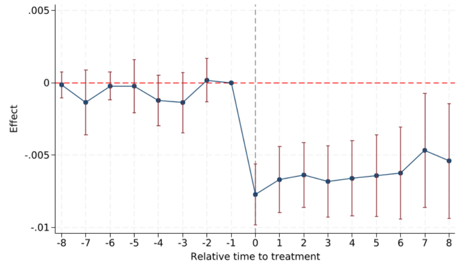

(d) Share paid apps

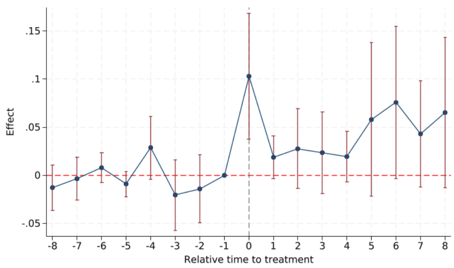

(e) ΔMonetiz. strategy

Figure 3.3: Comparative event study: Impact of M&A on app development (Part 2)

Notes: Continuation of the event study results from Figure 3.2. These figures display the impact of M&A on different app-related outcomes, following the methodology of Callaway and Sant'Anna (2021).

3.5.1.1 Mergers vs acquisitions

Table 3.7: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT (merger)	-0.0713	-0.121	0.0840	**1.089*****	-0.181	-0.263	0.094	0.008	-0.003
	(0.205)	(0.145)	(0.159)	(0.370)	(0.217)	(0.209)	(0.0775)	(0.00549)	(0.0204)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	6860	6860	6860	6860	6860	6860	6860	6860	6860
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (acquisition)	**-1.879***	**-1.682***	-0.059	**1.164****	**-0.877*****	**-0.557*****	**-0.320*****	**-0.006*****	**0.049****
	(1.002)	(0.999)	(0.096)	(0.562)	(0.220)	(0.189)	(0.104)	(0.002)	(0.0228)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	67857	67857	67857	67857	67857	67857	67857	67857	67857
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: This table presents the event study results segmented into two groups: M&A. This segmentation allows for a more detailed understanding of how different types of consolidation events impact app developer behaviour. The Average Treatment Effect on the Treated (ATT) is estimated separately for M&A, covering several key app-related outcomes.

The results in Table 3.7 reveal distinct post-transaction adjustment patterns following M&A. For mergers, the primary response occurs along the intensive margin. Relative to the counterfactual path absent the transaction, developers subject to a merger increase update activity by 1.089 updates per developer-period, indicating greater emphasis on maintaining and refining existing apps. In contrast, mergers do not generate statistically significant deviations from counterfactual trends in extensive-margin outcomes, including new app creation, app purchases, removals, dismissals, or sales. This suggests that mergers mainly alter how developers manage their existing portfolios rather than their scale or composition.

Acquisitions, by comparison, induce broader and more pronounced departures from counterfactual behaviour. Relative to what acquiring developers would have done absent the acquisition, they create 1.879 fewer new apps and purchase 1.682 fewer apps from other developers, reflecting a sharp contraction in expansionary activity. At the same time, acquisitions are associated with an increase of 1.164 updates per developer-period, consistent with intensified efforts to integrate, improve, and reposition retained apps. Portfolio turnover also declines following acquisitions: developers remove 0.877 fewer apps, dismiss 0.557 fewer apps, and sell 0.320 fewer apps than under the counterfactual, indicating a shift toward retention and internal consolidation rather than active reallocation through exit or ownership transfers.

Acquisitions further affect monetization behaviour. While the share of paid apps falls modestly relative to the counterfactual (-0.006), the likelihood of changing monetization strategy increases (+0.049), pointing to experimentation with pricing formats and revenue models. Taken together, these results indicate that while both M&A lead developers to optimize existing app portfolios, acquisitions trigger a more substantial reorientation away from growth through entry and acquisitions toward consolidation and internal restructuring.

Mergers | **Acquisitions**

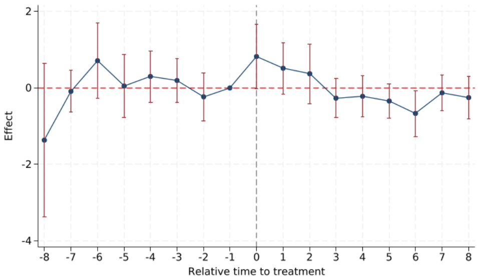

(a) #New apps

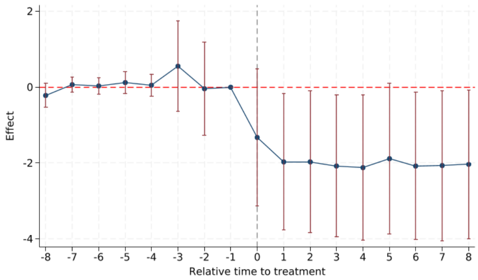

(b) #New apps

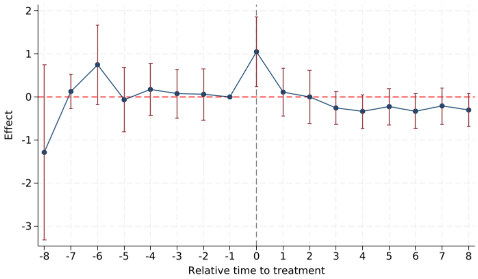

(c) #Purchased apps

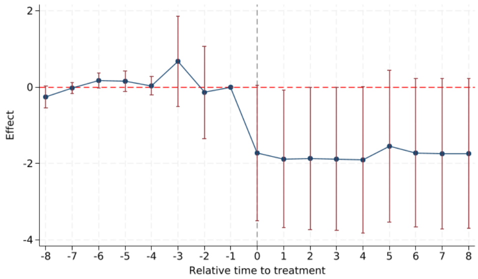

(d) #Purchased apps

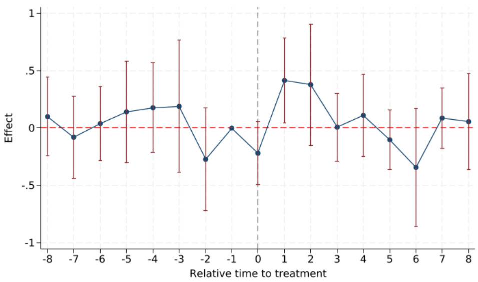

(e) #Newborn apps

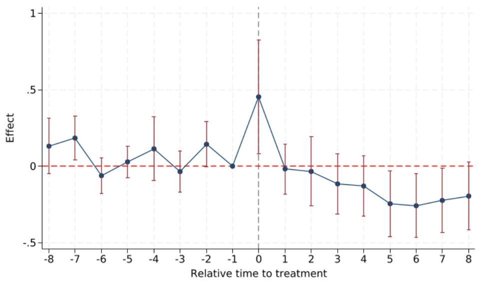

(f) #Newborn apps

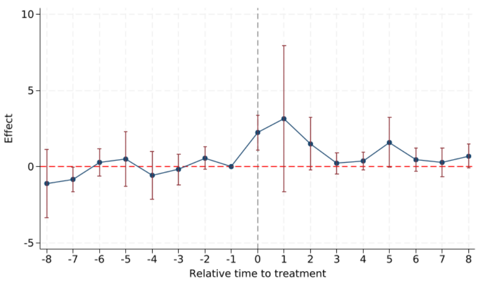

(g) #App updates

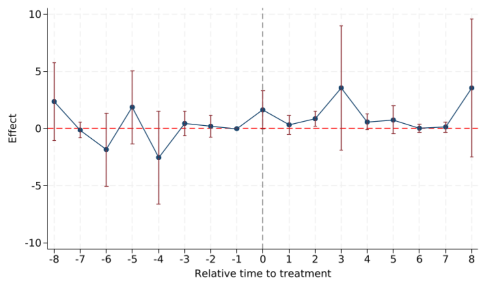

(h) #App updates

Figure 3.4: Comparative event study: Mergers vs Acquisitions and app development outcomes (Part 1)

Notes: This figure presents event-study estimates of the effect of M&A on developer outcomes, separately for mergers and acquisitions. The x-axis reports event time in six-week intervals, with 0 denoting the first post-event period.

Mergers **Acquisitions**

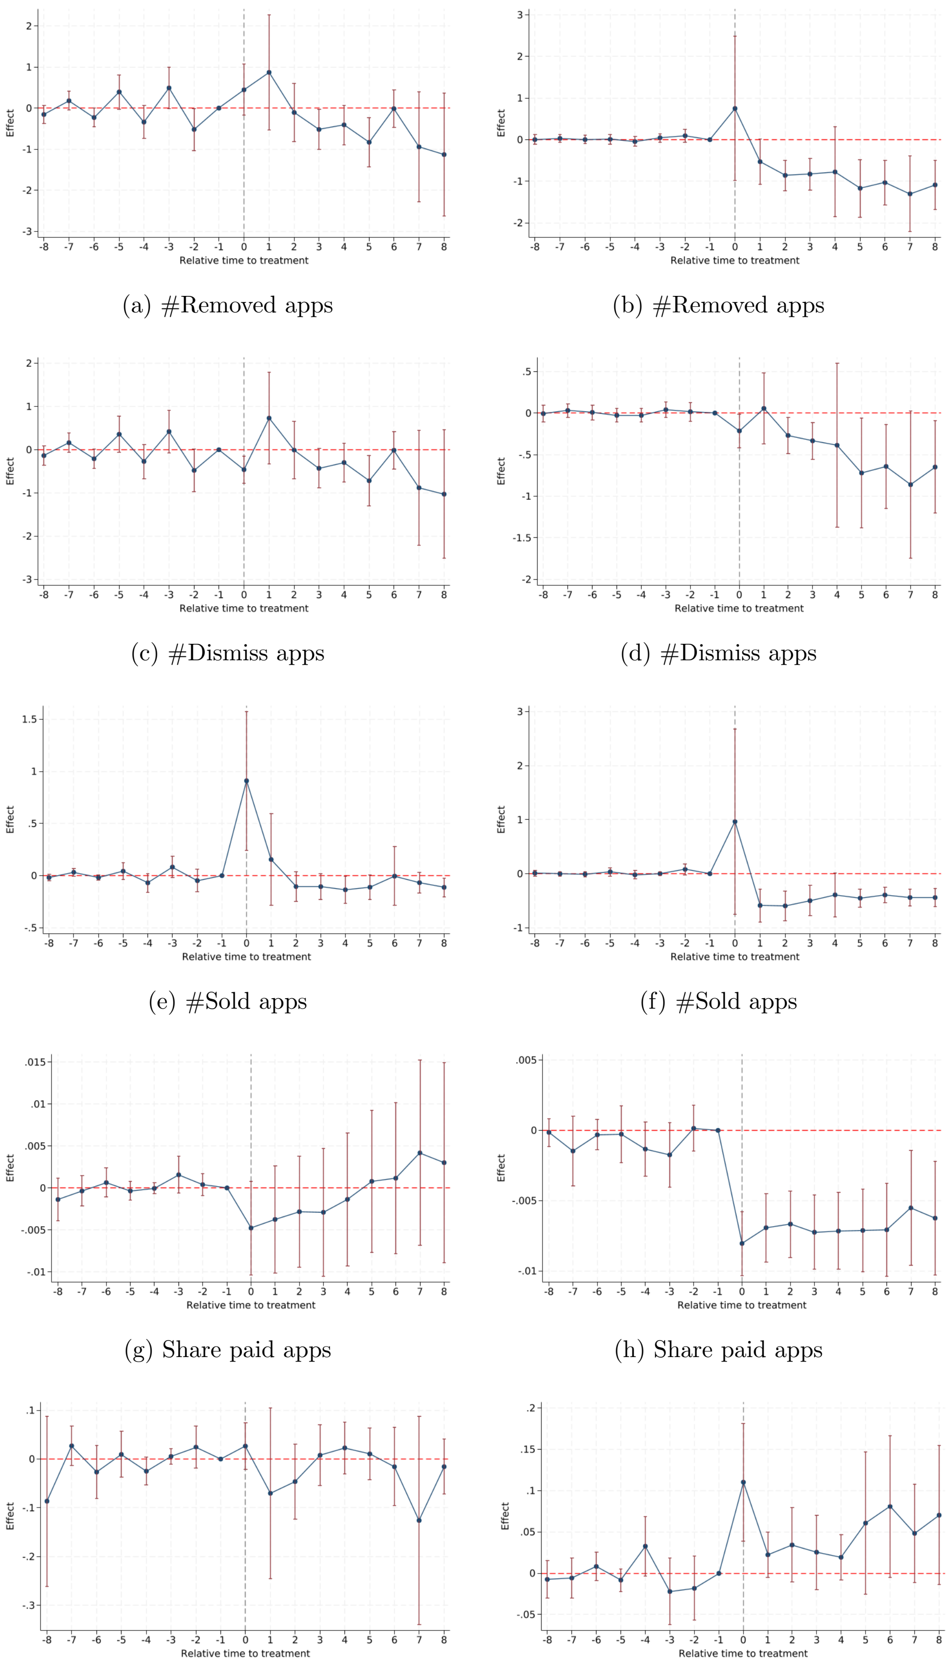

(a) #Removed apps (b) #Removed apps

(c) #Dismiss apps (d) #Dismiss apps

(e) #Sold apps (f) #Sold apps

(g) Share paid apps (h) Share paid apps

(i) ΔMonetiz. strategy (j) ΔMonetiz. strategy

Figure 3.5: Comparative event study: Mergers vs Acquisitions and app development outcomes (Part 2; continuation)

The event-study results in Figure 3.4 and Figure 3.5 highlight clear differences between mergers and acquisitions in how developers adjust their portfolios. Following mergers, expansionary margins—such as new app creation, purchased apps, and newborn apps-remain broadly stable, showing no systematic contraction. Instead, mergers are primarily associated with increased update activity and reductions in app removals, dismissals, and sales, consistent with a reallocation of effort toward maintaining and consolidating existing portfolios rather than expanding them.

Acquisitions display a markedly different pattern. Post-acquisition, developers substantially reduce new app creation and external app purchases, indicating a pronounced pull-back from expansionary activity. At the same time, update activity increases, and ownership transfers and app exits decline, pointing to active consolidation and within-portfolio reorganization following acquisitions.

Monetization responses further distinguish the two deal types. While mergers are associated with relatively limited changes in monetization, acquisitions coincide with a modest decline in the share of paid apps and greater adjustment in monetization strategies, suggesting increased experimentation with alternative revenue models. Overall, while both mergers and acquisitions lead to portfolio optimization, acquisitions induce sharper contractions in external growth and stronger reallocation across development and monetization margins.

In summary, mergers encourage a focus on portfolio optimization without significantly reducing expansionary activities such as new app creation, whereas acquisitions induce a more pronounced contraction in external growth, particularly through reductions in new app launches and app purchases. Both types of consolidation events are associated with increased update activity, indicating a reallocation of development effort toward existing apps. However, the strategic emphasis differs: mergers are associated with greater portfolio stability, while acquisitions lead to more substantial shifts away from expansionary strategies and toward within-portfolio consolidation.

3.5.1.2 Horizontal versus non-horizontal M&A

Table 3.8 documents systematic differences in post-M&A adjustment between horizontal and non-horizontal transactions. For non-horizontal M&A, developers exhibit clear and statistically significant deviations from the counterfactual path in which no transaction occurred. Relative to this no-M&A counterfactual, new app creation is lower by 0.666 and app purchases are lower by 0.522, indicating a contraction in expansionary activity following non-horizontal consolidation.

At the same time, non-horizontal M&A are associated with a significant reallocation of effort toward existing products. Relative to the counterfactual, update activity increases by 0.986, while portfolio turnover declines markedly: removals fall by 0.903, dismissals by 0.603, and app sales by 0.299. Together, these effects indicate that non-horizontal transactions lead developers to stabilize and consolidate their portfolios, reducing both entry and exit while intensifying management of incumbent apps.

By contrast, horizontal M&A display fewer statistically robust departures from the counterfactual. The primary significant adjustment occurs along the intensive margin, with update activity higher by 1.406 relative to the no-M&A counterfactual. Other margins do not exhibit systematic or precisely estimated changes. This pattern suggests that, when developers overlap in product space prior to the transaction, post-M&A adjustment operates mainly through changes in how existing apps are maintained rather than through systematic changes in portfolio size or composition.

Overall, the results imply that consolidation following M&A is most clearly and systematically expressed in non-horizontal transactions, where developers scale back expansionary activity and reduce portfolio churn relative to the counterfactual path, whereas horizontal transactions primarily affect intensive-margin behaviour.

Table 3.8: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT (Horizontal M&A)	-3.521	-3.352	0.102	**1.406****	**-0.738****	**-0.470****	-0.268	-0.004	0.0749
	(2.483)	(2.480)	(0.234)	(0.645)	(0.349)	(0.236)	(0.199)	(0.004)	(0.048)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	27858	27858	27858	27858	27858	27858	27858	27858	27858
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (Non-horizontal M&A)	**-0.666*****	**-0.522*****	-0.095	0.986	**-0.903*****	**-0.603*****	**-0.299*****	**-0.006*****	**0.031***
	(0.166)	(0.129)	(0.090)	(0.721)	(0.201)	(0.174)	(0.097)	(0.002)	(0.019)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	48039	48039	48039	48039	48039	48039	48039	48039	48039
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: This table presents the ATT results for horizontal and non-horizontal M&A. The results are segmented into four groups: horizontal mergers, non-horizontal mergers, horizontal acquisitions, and non-horizontal acquisitions. The table examines the impact of these events on various app-related metrics, such as new apps, purchased apps, newborn apps, app updates, and more. Row definitions refer to Table 3.6.

Horizontal **Non-horizontal**

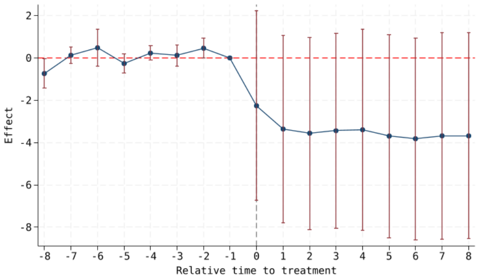

(a) #New apps

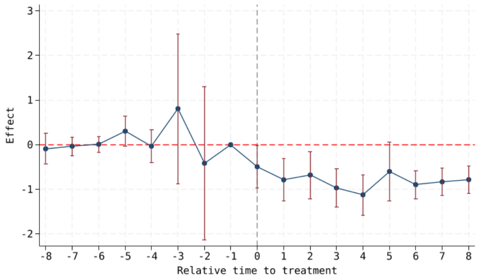

(b) #New apps

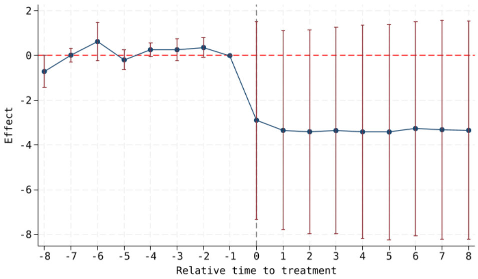

(c) #Purchased apps

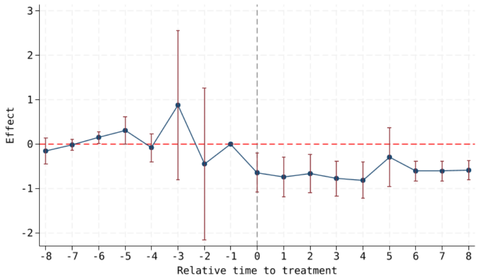

(d) #Purchased apps

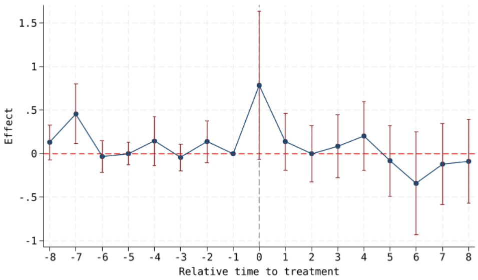

(e) #Newborn apps

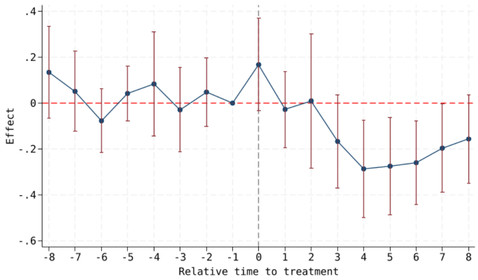

(f) #Newborn apps

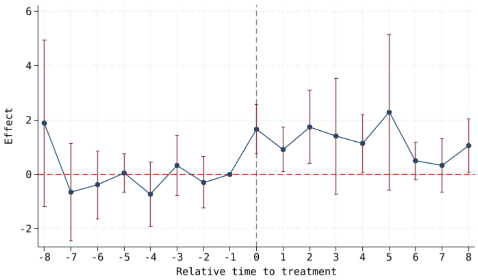

(g) #App updates

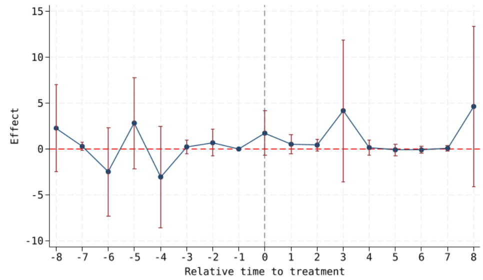

(h) #App updates

Figure 3.6: Comparative event study: Horizontal vs Non-horizontal M&A and app development outcomes (Part 1)

Notes: This figure (Parts 1–2) presents event-study estimates of the effect of M&A on developer outcomes, separately for horizontal and non-horizontal transactions, following Callaway and Sant'Anna (2021). The x-axis reports event time in six-week intervals, with 0 denoting the first post-event period.

Horizontal **Non-horizontal**

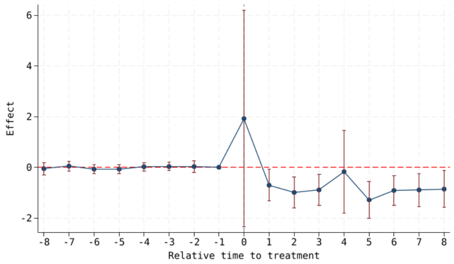

(a) #Removed apps

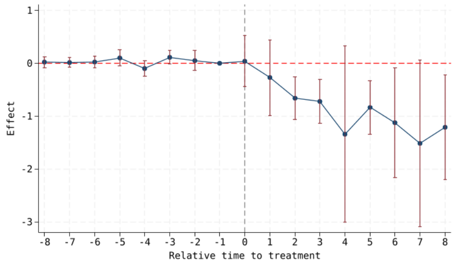

(b) #Removed apps

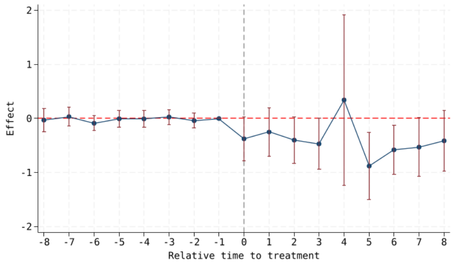

(c) #Dismiss apps

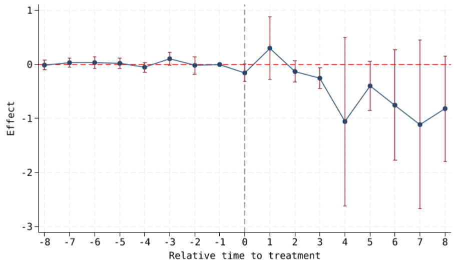

(d) #Dismiss apps

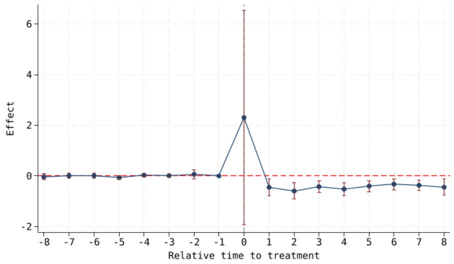

(e) #Sold apps

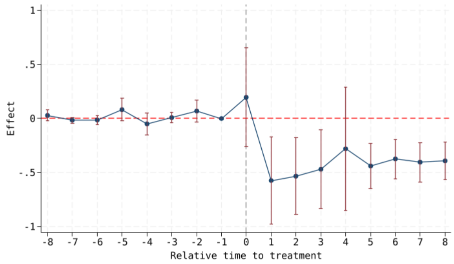

(f) #Sold apps

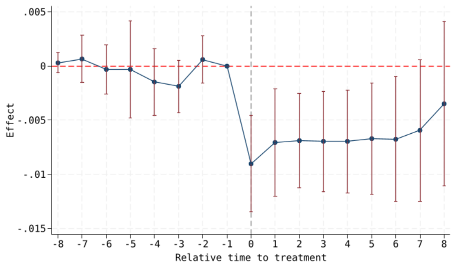

(g) Share paid apps

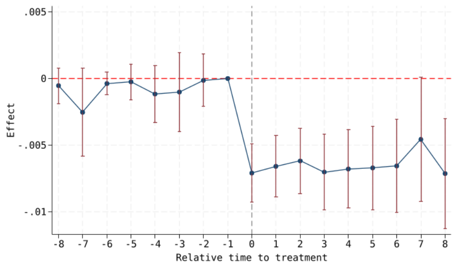

(h) Share paid apps

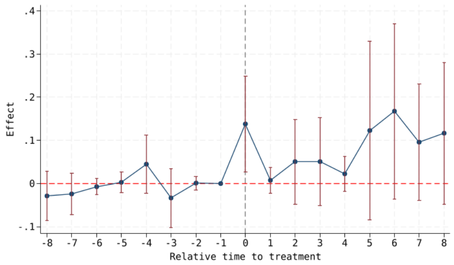

(i) ΔMonetiz. strategy

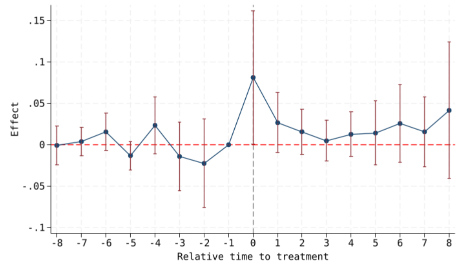

(j) ΔMonetiz. strategy

Figure 3.7: Comparative event study: Horizontal vs Non-horizontal M&A and app development outcomes (Part 2)

The event-study figures (Figure 3.6 and Figure 3.7) further illustrate these patterns by contrasting horizontal and non-horizontal M&A within a unified framework. For horizontal M&A (left-hand panels), the estimates show pronounced reductions in new app creation and purchased apps, alongside modest but persistent increases in app updates. These dynamics point to a shift away from expansionary activity toward within-portfolio maintenance and consolidation. Consistent with this interpretation, horizontal transactions are also associated with fewer removed, dismissed, and sold apps, indicating greater portfolio stability following consolidation.

For non-horizontal M&A (right-hand panels), the contraction in new app creation and external acquisitions is generally less pronounced and, in some cases, partially reverses over time. Effects on update activity and portfolio adjustment margins are more moderate, with smaller declines in removals and dismissals relative to horizontal deals. Nonetheless, both types of M&A exhibit a sharp and persistent decline in app sales, suggesting a common post-merger strategy of retaining assets within the developer's portfolio.

Overall, the comparison highlights that horizontal M&A exert stronger negative effects on expansionary margins—such as new app creation and acquisitions-while non-horizontal M&A display milder and more heterogeneous responses across app development outcomes.

3.5.2 Placebo test

To confirm the validity of these results, we conducted placebo tests by setting false treatment periods one to six periods before the actual M&A. The results, shown in Figure 3.10 in the appendix, consistently fluctuate around zero, confirming that the observed effects are not driven by random noise or pre-existing trends.

3.5.3 Robustness checks

In addition to placebo tests, we conducted several robustness checks.

First versus all M&A events: We performed a robustness check by varying the event definition, as shown in Table 3.14. In the main analysis, we included only the first M&A event for each developer to ensure observation independence. As a robustness test, we also examined an alternative specification that includes all M&A

events experienced by each developer during the sample period. Across both definitions, the findings are highly consistent: M&A are associated with significant reductions in new app creation and purchased apps, as well as significant increases in app updates. Similarly, there are robust declines in app removals, dismissals, and sales, alongside a slight reduction in the share of paid apps and an increase in monetization experimentation. The close alignment of results across event definitions confirms the generalizability and persistence of the observed effects, reinforcing the robustness of our main conclusions.

Different pre-treatment periods: To further test the robustness of our findings, we conducted additional analyses using different pre-treatment periods. Specifically, we examined the results for 4, 6, and 8 pre-treatment periods to assess the sensitivity of our estimates as shown in Table 3.15. The 8-period analysis serves as our main result, and the results from the 6- and 4-period analyses exhibit patterns consistent with the main findings. While slight variations are observed, the overall trends remain stable: M&A are associated with reductions in new app launches and app variety, as well as increased app updates and shifts in monetization strategies. These results reinforce the robustness of our findings, demonstrating that the observed effects persist regardless of the specific pre-treatment period chosen for analysis. The alignment across different pre-treatment periods underscores the reliability of our conclusions and the robustness of our empirical approach.

Japanese Google Play data: To ensure the reliability and generalizability of our findings, we conducted a robustness check using data from the Japanese version of the Google app market. The Japanese dataset was available quarterly, compared to the six-weekly intervals in our primary analysis of the US Google Play data. As we did for the US Google Play, we identified all M&A that occurred in the Japanese Google Play app market. When comparing the entire population data across the US (Table 3.1) and Japan (Table 3.19), the summary statistics reveal similarities. Both datasets show that M&A are consistently associated with fewer new app launches and reduced app variety, as well as significant differences in monetization strategies when compared to non-M&A developers. These comparisons reinforce the generalizability of our findings across diverse markets and datasets. Additionally, we compared the empirical event study ATT results from the Japanese quarterly data with those from the US quarterly data to ensure consistency across temporal resolutions as shown in Table 3.20. The quarterly US data reveal patterns highly similar to those observed in the six-weekly analysis, further supporting the robustness of our findings. The alignment between the Japanese and US quarterly results underscores that the effects of M&A persist across different market contexts and temporal granularities, reinforcing the reliability and generalizability of our conclusions.

3.5.4 Main results at app-level results

This section examines the impact of M&A on app-level outcomes. Relative to the developer-level analysis, this approach exploits a larger number of observations but focuses on a narrower set of behavioural margins. Matching is performed at the app level, allowing for a more granular assessment of changes in update activity, portfolio status (dismissal and sale), pricing, and monetization choices following M&A.

Table 3.9: Event study ATT results

	App update	Dismissed following period	Sold following period	Price ($)
ATT	**0.042*****	**0.039*****	**-0.260*****	**0.111***
	(0.043)	(0.001)	(0.001)	(0.017)
Time FE	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓
Observations	791047	777590	750015	791047
PT†	YES	YES	NO	YES

Notes: The table reports ATT estimates from the Callaway and Sant'Anna (2021) estimator, examining the impact of M&A on various app-level outcomes. Standard errors (in parentheses) are clustered by app ID. All models include time and app fixed effects. The price specification is estimated on the paid-app subsample (Price¿0); for the full sample, most apps have zero posted price. PT† indicates that the parallel trends assumption was tested and holds. Results shown use 8 pre-treatment periods, matching units accordingly. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.9 reports average treatment effects at the app level, focusing on short-run changes in app maintenance, exit, and pricing behaviour following a developer merger or acquisition.

Consistent with the developer-level evidence, M&A significantly increase the probability that an app receives an update in the subsequent period. The estimated ATT of 0.042 implies a 4.2 percentage point increase in the likelihood of an update relative to a pre-treatment baseline of 0.16, corresponding to an increase of approximately 26%. This finding indicates a shift towards greater maintenance and incremental improvement of existing apps after consolidation.

At the same time, M&A affect app exit margins in economically important ways. The probability that an app is dismissed from the platform increases by 3.9 percentage points in the period following an acquisition, relative to a baseline dismissal probability of 3.6%, implying a substantial rise in exit risk. By contrast, the likelihood that an app is sold in the subsequent period declines sharply after M&A. Because app sales are extremely rare events (mean ≈ 0.018) and the estimates are obtained from a linear DiD framework, the magnitude of this effect should not be interpreted as a literal probability change. Instead, the result indicates a strong reduction in secondary-market ownership transfers following consolidation. Taken together, these findings suggest that post-merger adjustment operates primarily through increased app pruning rather than through market-based reallocation of ownership.

Turning to pricing outcomes, Table 3.9 shows a positive and statistically significant effect on app prices among paid apps. The estimated ATT of 0.111 corresponds to an average price increase of approximately 11% relative to a pre-treatment mean price of 0.20. Because the vast majority of apps on Google Play are offered for free, this effect is identified from a small subset of the app ecosystem and should not be interpreted as a platform-wide increase in consumer prices. Accordingly, while statistically precise, the aggregate economic impact of the price response is limited. By contrast, adjustments along non-price margins—such as update intensity, app dismissals, and portfolio restructuring-are larger in relative terms and affect a substantially broader set of apps.

Taken together, these results indicate that developer M&A are not strictly price-neutral for paid apps, but that their primary competitive effects operate through non-price channels rather than through widespread changes in consumer prices.

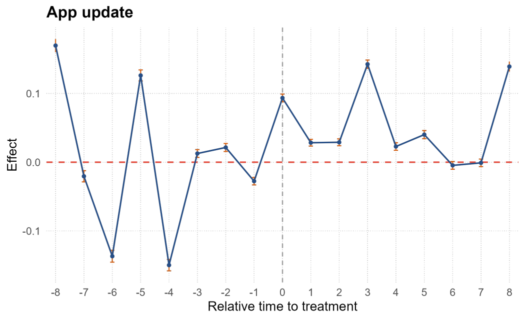

(a) App update

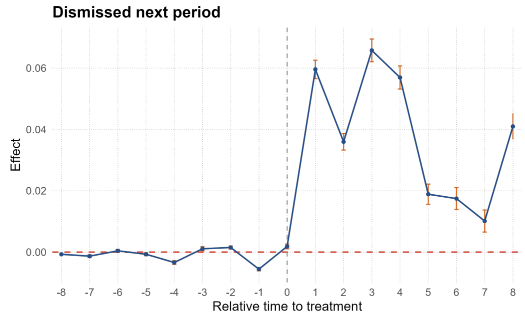

(b) Dismissed following period

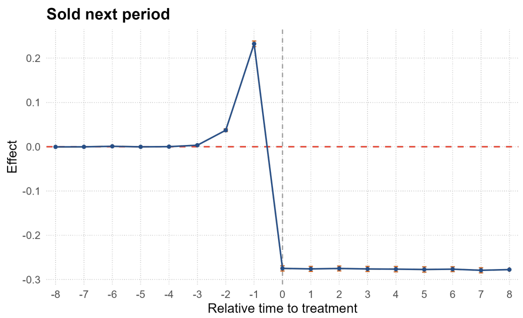

(c) Sold following period

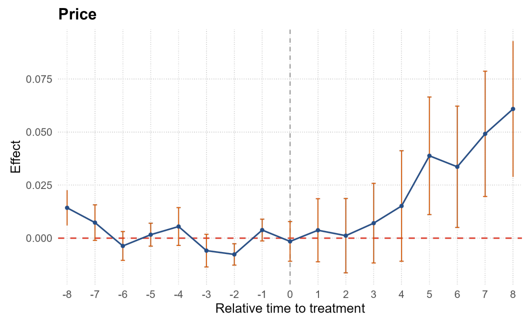

(d) Price($)

Figure 3.8: Comparative event study: Impact of M&A on app development

Notes: These figures show the event study results for mergers, capturing how various app-related metrics evolve before and after the merger event. The x-axis represents date relative to the merger, with the six week interval after the merger set at 0. Each sub-figure reflects a different app-related outcome, providing insights into the short- and long-term effects of mergers on app developer behaviour.

As shown in Figure 3.8, the event–study analysis reveals distinct and economically important post-merger dynamics in app-level outcomes. Pre-treatment coefficients remain close to zero across all panels, providing reassuring support for the parallel-trends assumption.

Immediately following the M&A event, the probability that an app receives an *update* rises sharply. This increase persists for several periods, indicating that consolidation is associated with intensified maintenance and incremental improvement of existing apps rather than with continued entry. The dynamic pattern mirrors the positive average treatment effect reported in Table 3.9 and supports the interpretation that post-merger developer effort is reallocated toward within-portfolio optimization.

At the same time, the probability that an app is *dismissed* increases in the short run, consistent with selective pruning of underperforming titles, while the probability that an app is *sold to another developer* declines sharply after the merger. This collapse in app sales suggests that, following consolidation, developers are more likely to retain and internally reorganize their app portfolios rather than transfer assets across firms.

Finally, app *prices* display a gradual upward adjustment beginning a few periods after treatment. This pattern is driven by paid apps only and is modest in absolute magnitude, but statistically precise, consistent with the positive price effect reported in Table 3.9. Given that the vast majority of apps are offered for free, these results should not be interpreted as a broad increase in consumer prices, but rather as limited monetization adjustments within the paid-app segment once post-merger integration stabilizes.

Taken together, the event–study evidence reinforces the conclusion that developer M&A in the app economy primarily operate through non-price channels. Post-merger adjustments are concentrated on maintenance intensity, portfolio consolidation, and internal reallocation of development effort, with price changes playing a secondary and economically limited role.

3.6 Conclusion

This study analyses 126 mergers and 1,246 acquisitions in the mobile app market between 2016 and 2017, using data from 8 million apps and 1.8 million developers on the US Google Play Store (2015–2019). Our findings show that M&A generally lead to a reduction in new app launches and external app acquisitions, alongside an increase in app updates. This pattern indicates a strategic shift among developers toward consolidating and intensifying effort within existing product portfolios rather than pursuing expansion through new app creation or acquisitions.

When comparing M&A, acquisitions have a stronger negative impact on new app creation, with developers scaling back expansionary activities post-acquisition. Mergers, in contrast, are associated with particularly large increases in update activity, consistent with greater emphasis on integration and ongoing development within the merged portfolio. Further segmentation reveals distinct patterns for horizontal and non-horizontal M&A. At the developer level, portfolio reshuffling manifests as fewer ownership transfers (sales) and changes in aggregate removal components, while at the app level we observe selective pruning-an elevated probability that an incumbent title is dismissed in the subsequent period-alongside a sharp decline in app sales.

M&A also influence monetization strategies at the developer level. Developers adjust monetization choices following consolidation, as reflected in changes in pricing models and the allocation between free and paid apps. Although the share of paid apps declines slightly, this points to a reorientation toward alternative monetization mechanisms rather than price-based competition. These findings underscore the importance of non-price margins in digital markets, where competitive adjustment often operates through development intensity, portfolio composition, and monetization design rather than prices alone.

From an app-level perspective, M&A are associated with more frequent maintenance activity (updates), more active pruning (dismissals), and a sharp decline in ownership transfers (sales), consistent with post-merger within-portfolio reallocation. Prices increase modestly within the paid-app segment (Table 3.9), but because the vast majority of apps are offered for free, this implies only a limited aggregate change in consumer prices and does not alter the central conclusion that post-M&A adjustment in this ecosystem operates primarily through non-price margins. These conclusions are robust to replicating the design in the Japanese Google Play market using quarterly data: despite lower precision due to coarser temporal aggregation, the estimates show the same qualitative reallocation pattern-less external expansion (notably fewer purchased apps) and a lower share of paid apps, alongside broadly similar movements in update and portfolio-adjustment margins-suggesting that the core non-price mechanisms generalize across national markets.

Finally, this study highlights several directions for future research. Measuring quality outcomes directly—such as user ratings, engagement, or retention-would allow a sharper assessment of how changes in development intensity translate into consumer welfare. More broadly, defining economically meaningful markets in digital settings remains challenging due to limited price variation, and further work is needed to understand how M&A affect developers of different sizes and competitive positions within platform ecosystems.

3.A Appendix

3.A.1 Regulatory actions in digital markets targeting app developers

The increasing dominance of tech giants in digital markets has led to heightened regulatory scrutiny of M&A involving app developers as shown in Table 3.10. While many regulatory actions focus on market share, revenue, and pricing strategies, they also target broader issues such as innovation, competition, developer autonomy, and user rights. Regulatory bodies, such as the European Commission and the Japan Fair Trade Commission, have taken significant action against companies like Apple and Google for practices affecting app developers. These cases address issues ranging from restrictive platform rules to unfair contractual terms and barriers to alternative payment methods, impacting both price and non-price dimensions of competition.

3.A.2 GAFAM mergers and acquisitions (2016-2017)

This section presents an overview of M&A M&A from 2016-2017, highlighting their impact on app-related outcomes. The data, manually collected and cross-verified, as shown in Table Table 3.11. In our main results, we apply the 'excluding 2015 M&A' criterion, which removes developers involved in M&A events during 2015. As a result, developers that retain M&A-associated names post-acquisition, such as Microsoft Corporation's acquisition of AltspaceVR, are excluded from the main sample. This exclusion simplifies the matching process since it is challenging to find comparable control developers for these major tech firms. By removing such entities, we enhance the robustness of our results, ensuring that comparisons are made between more similar treatment and control units, thereby reducing potential bias introduced by unmatched M&A-associated developers.

The event study results in Table 3.12 analyse the ATT for various app-related outcomes following GAFAM acquisitions. Notably, there is a statistically significant increase in the number of newborn apps (ATT = 0.212), suggesting a focus on launching new products or features post-acquisition. Other metrics, such as new apps, app updates, and monetization strategies, show no significant short-term changes, indicating that M&A acquisitions primarily target innovation and new app

Table 3.10: Regulatory actions addressing app developer ecosystems

Year	**Case**	**Effect**	**Regulatory Body**
2021	JFTC closes App Store investigation	Apple allowed developers of reader apps to include external links for purchases, reducing Apple's control over in-app purchases.	Japan Fair Trade Commission
2021	ACM v Apple	Apple required to allow alternative payment methods for dating apps, reducing its dominance over app transactions.	Netherlands Authority for Consumers and Markets (ACM)
2022	CMA investigation into Apple	The UK CMA concluded its investigation into Apple's mobile ecosystems, opting to rely on the new regulatory regime under the Digital Markets, Competition and Consumers Act.	UK CMA
2022	Ministère de l'économie et des finances v Google	Google was found to have imposed unfair terms on developers, including unilateral contract modifications and restrictive pricing conditions, limiting developers' autonomy.	Tribunal de Commerce de Paris
2023	Bundeskartellamt and Apple	Apple was designated as a company of 'paramount significance' across markets, placing it under closer scrutiny but without immediate penalties or conduct changes.	Bundeskartellamt (Germany)
2023	AGCM v Apple	AGCM investigated Apple for allegedly discriminatory privacy policies that required stricter consent prompts for third-party app tracking compared to Apple's own, favoring Apple's ecosystem and potentially limiting competition.	Italian Competition Authority (AGCM)
2024	European Commission (Apple App Store)	Ongoing investigation into Apple's restrictions on third-party payment systems, potentially reshaping App Store rules across Europe.	European Commission

Table 3.11: GAFAM Mergers and acquisitions (2016-2017)

Target Entity	**Name after Acquisition**	**Acquirer (GAFAM)**	**Event Date**
LinkedIn	LinkedIn	Microsoft Corporation	June 13, 2016
Xamarin Inc.	Xamarin Inc.	Microsoft Corporation	March 17, 2016
International Software	International Software	Microsoft Corporation	April 18, 2017
AltspaceVR	Microsoft Corporation	Microsoft Corporation	October 3, 2017
Beam Interactive, Inc	Microsoft Corporation	Microsoft Corporation	August 8, 2016
Moodstocks	Moodstocks	ALPHABET	July 6, 2016
Anvato Inc	Anvato Inc	ALPHABET	July 8, 2016
Cloud9 Inc.	Cloud9 Inc.	AMAZON	July 14, 2016
Apigee Admin	Apigee Corporation	ALPHABET	September 8, 2016
Apigee Corporation	Apigee Corporation	ALPHABET	September 8, 2016
TwitterDev	Google Fabric	ALPHABET	January 19, 2017
WORKFLOW	WORKFLOW	Apple	March 23, 2017
Beddit	Beddit	Apple	May 8, 2017
Masquerade Technologies, Inc	Facebook	Facebook	March 9, 2016
Nascent Objects Inc.	Nascent Objects Inc.	META	September 19, 2016
FacioMetrics LLC	FacioMetrics LLC	META	November 16, 2016
COLIS PRIVE	COLIS PRIVE	AMAZON	January 11, 2016
Curse, Inc.	Curse, Inc.	AMAZON	August 16, 2016
Thinkbox Software Inc.	Thinkbox Software Inc.	AMAZON	March 6, 2017
Whole Foods Market, Inc.	Whole Foods Market, Inc.	AMAZON	June 16, 2017
Owlchemy Labs	Owlchemy Labs	ALPHABET	May 10, 2017
Souq.com	Souq.com	AMAZON	July 3, 2017

Notes: This table presents the M&A conducted by GAFAM (Google, Amazon, Facebook, Apple, and Microsoft) between 2016 and 2017. The **Target Entity** column lists the acquired companies, while the **Name after Acquisition** column displays the company before the M&A. The **Acquirer (GAFAM)** column shows which of the five tech giants was the acquirer, and the **Event Date** indicates date which the M&A took place.

Table 3.12: Event study ATT results (GAFAM)

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT(Frist quantile)	0.126	-0.131	0.212**	-0.199	0.481	0.483	0.333	0.00000421	-0.200
	(0.114)	(0.171)	(0.107)	(0.470)	(0.345)	(0.352)	(0.272)	(0.0000965)	(0.193)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	952	952	952	952	952	952	952	952	952
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: The table reports the ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8 pre-treatment periods, where the matching is applied to units with at least 8 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

creation rather than broad portfolio restructuring. In summary, M&A M&A demonstrate a strategic emphasis on fostering innovation through new app launches while maintaining stability in other aspects of app management and monetization. This targeted approach underscores their focus on enhancing product offerings rather than immediate, large-scale portfolio changes.

3.A.3 M&A definition

Merger: There are two app developers (D and E) in T. Then D and E merged into one new app developer F. The app developer F should be a brand new app developer by checking the ownership situation of app developer F. Additionally, app developers (D and E) should not have apps after T+1. This case can be called a merger.
Acquisition: There are three app developers (A, B, and C) in T. Then A bought B and C. App developer A should have apps both before and after T+1 by checking the ownership situation of app developer A. Additionally, app developers (B and C) should not have apps after T+1. This case can be called an acquisition.

Table 3.13: App ownership and transactions example before and after the acquisition

Time Period	Developer Entity	Apps Owned	Newborn Apps	Purchased Apps	New Apps	Sold Apps	Dismissed Apps	Removed Apps
$t-2$	Developer A	a, b	None	None	0	None	None	0
	Developer B	e	None	None	0	None	None	0
	Developer C	x, y, h	None	None	0	None	None	0
	Developer D	p, q, z	None	None	0	None	None	0
$t-1$	Developer A	a, b, c, d	c, d	None	2	None	None	0
	Developer B	e, f, g, h	f, g	h (from C)	3	None	None	0
	Developer C	x, y	None	None	0	h (to B)	None	1
	Developer D	p, q, z, r	r	None	1	None	None	0
t	A + B (Acquisition Event)	a, b, c, d, g, h, z	None	z (from D)	1	e (to C)	f	2
	Developer C	x, e	None	e (from A + B)	1	None	y	1
	Developer D	p, q, r, s	s	None	1	z (to A+B)	None	0
$t+1$	A + B (Post-acquisition)	a, c, d, g, h, i	i	None	1	b, f	None	2
	Developer C	x, e, y	None	None	0	None	None	0
	Developer D	p, q, r, s, t	t	None	1	None	r	1
$t+2$	A + B (Post-acquisition)	c, d, g, h, i, j	j	None	1	e	None	1
	Developer C	a, x, y, e	None	a (from A + B)	1	None	None	0
	Developer D	p, s, t, u	u	None	1	None	q	1

Notes: This table provides a detailed overview of app ownership and app transactions among Developers A, B, C, and D across multiple time periods surrounding the acquisition of Developer B by Developer A. The table tracks newly developed apps (Newborn Apps), apps purchased from other developers (Purchased Apps), and the total number of new apps in each period (New Apps). It also records apps sold to other developers (Sold Apps), apps dismissed from the market (Dismissed Apps), and the total number of removed apps in each period (Removed Apps). Each developer's app portfolio is updated accordingly across time periods $t-2$ through $t+2$, highlighting the dynamics of the app market before, during, and after the acquisition event.

Table 3.13 illustrates the process of app acquisition and removal during an event study involving the merger of Developer A and Developer B at time t, along with interactions involving Developer C and D. It captures the changes in the app portfolios of the developers over different time periods, tracking both the new and removed apps. The breakdown of apps includes: Newborn Apps: Apps developed internally by the developer during each period. Purchased Apps: Apps acquired from other developers. Sold Apps: Apps sold to other developers. Dismissed Apps: Apps completely removed from the market by the developer. The table includes two additional columns, Number of New Apps and Number of Removed Apps, which summarize the total number of new and removed apps in each period by summing the respective components (newborn + purchased for new apps, and sold + dismissed for removed apps).

3.A.4 Alternative count of merger and acquisition

In this section, we will count each merger and acqusition into the sample.

Additionally, we will explore an alternative method that accounts for multiple M&A events. For example, if developer A acquires B at time t and then acquires C at time t+3, we will treat A+B as a combined entity until time t+3, when the second acquisition occurs. At that point, we will redefine the entity as A+B+C to account for the new event. This approach allows us to incorporate subsequent acquisitions into the analysis.

Table 3.14: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT (First event)	**-1.723***	**-1.552***	-0.045	**1.145****	**-0.789*****	**-0.504*****	**-0.285*****	**-0.005*****	**0.046****
	(0.918)	(0.914)	(0.0888)	(0.514)	(0.196)	(0.166)	(0.0947)	(0.00196)	(0.0211)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	73876	73876	73876	73876	73876	73876	73876	73876	
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (All events)	**-1.741***	**-1.561***	-0.054	**1.106****	**-0.796*****	**-0.512*****	**-0.284*****	**-0.005****	**0.046****
	(0.915)	(0.912)	(0.0897)	(0.514)	(0.198)	(0.168)	(0.0947)	(0.00193)	(0.0211)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	74153	74153	74153	74153	74153	74153	74153	74153	74153
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: This table presents robustness checks for the ATT results using different event definitions. The first panel (**ATT (First event)**) includes only the first M&A event for each developer, following the main analysis. The second panel (**ATT (All events)**) includes every M&A event observed for each developer during the sample period, capturing repeated M&A activities. Standard errors are clustered by app developer ID. PT† indicates whether the parallel trend assumption holds. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results in Table 3.14 provide a robustness check for the main findings by comparing the impact of M&A when including only the first observed event per developer (**ATT (First event)**) versus considering all M&A events experienced by each developer (**ATT (All events)**). Across both specifications, the results remain highly consistent: M&A are associated with significant reductions in new app creation (-1.72 to -1.74) and purchased apps (-1.55 to -1.56), as well as marked increases in app updates (1.11 to 1.15). Similarly, there are significant declines in removed apps, dismissed apps, and sold apps, reflecting a robust pattern of post-M&A portfolio consolidation. The share of paid apps also decreases slightly, while developers continue to experiment with monetization strategies. Overall, these robustness checks confirm that the direction, magnitude, and significance of the main effects are stable across alternative event definitions, further strengthening confidence in the validity and generalizability of the main results.

3.A.5 Using varying pre-treatment periods

This section presents the results of ATT for M&A on various app-related outcomes, with robustness checks using different pre-treatment periods during the matching process as shown in Table 3.15. The table reports results for 4, 6 and 8 pre-treatment periods, ensuring that the matching is applied to units with at least the corresponding number of pre-treatment observations. The purpose of this approach is to assess whether the observed effects are consistent across different pre-treatment lengths, thereby reinforcing the robustness and stability of the findings.

Overall, the results indicate that M&A have a consistent and significant impact on key outcomes such as the reduction in the number of new apps, purchased apps, and app removals, as well as an increase in app updates. These effects hold across different pre-treatment periods, with slight variations in magnitude. By controlling for various pre-treatment lengths, this analysis demonstrates that the observed M&A impacts are not driven by short-term fluctuations or pre-existing trends, thereby enhancing the credibility and generalizability of the main findings. This robustness check ensures that the effects of M&A on app development and portfolio management remain robust to different historical baselines.

Table 3.15: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT (8 periods)	**-1.723***	**-1.552***	-0.045	**1.145****	**-0.789*****	**-0.504*****	**-0.285*****	**-0.005*****	**0.046****
	(0.918)	(0.914)	(0.0888)	(0.514)	(0.196)	(0.166)	(0.0947)	(0.00196)	(0.0211)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	73876	73876	73876	73876	73876	73876	73876	73876	
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (6 periods)	**-1.600***	**-1.500***	0.036	**1.119***	**-1.811*****	**-1.462****	**-0.349*****	**-0.005****	-0.003
	(0.830)	(0.848)	(0.128)	(0.657)	(0.642)	(0.637)	(0.117)	(0.002)	(0.031)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	80134	80134	80134	80134	80134	80134	80134	80134	80134
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (4 periods)	**-1.481***	**-1.573****	0.202	**0.906***	**-0.653*****	**-0.240***	**-0.414*****	**-0.004****	0.0468**
	(0.791)	(0.791)	(0.164)	(0.499)	(0.190)	(0.140)	(0.111)	(0.002)	(0.022)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	84801	84801	84801	84801	84801	84801	84801	84801	84801
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

*Notes:*The table reports ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8, 6, and 4 pre-treatment periods, where the matching is applied to units with at least 8, 6, and 4 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.A.6 Comprehensive balance check for matched treatment and control groups

Table 3.16 and Table 3.17 present the results of a comprehensive balance check conducted after matching treated and control units using a 1:1 matching ratio with Mahalanobis distance. The analysis compares control and treatment groups across various outcome and control variables over several lead periods (from lead 1 to lead 8) before M&A events. The first table focuses on app-related outcome variables such as the number of new apps, purchased apps, app updates, removed apps, dismissed apps, and sold apps, while the second table reports control variables, including log-transformed measures of owned apps, ratings, average rating, paid apps, translated descriptions, and missing downloads.

The reported mean differences and p-values assess whether matching achieves balance across key variables between treatment and control groups. Significant differences (p-values less than 0.05) are highlighted, indicating areas where balance may be more difficult to achieve. Despite some variables showing persistent imbalances, particularly for outcome variables like the number of new apps, app updates and removed apps in specific lead periods, the overall results demonstrate that the matching procedure effectively minimizes differences between groups across most variables. This thorough balance check enhances the validity of subsequent

Table 3.16: Balance check after matching

PSM(1:1)	Control	Treatment	Mean-diff	P-value	N
Outcomes					
#New apps (lead 1)	1.15	2.66	1.508	0.07	2697
#New apps (lead 2)	0.75	1.22	0.465	**0.01**	2697
#New apps (lead 3)	0.57	1.38	0.809	0.16	2697
#New apps (lead 4)	0.52	0.93	0.411	**0.01**	2697
#New apps (lead 5)	0.68	0.97	0.295	0.15	2697
#New apps (lead 6)	0.57	0.98	0.402	**0.01**	2697
#New apps (lead 7)	0.66	0.98	0.316	0.10	2697
#New apps (lead 8)	0.47	1.01	0.539	**0.00**	2697
#Purchased apps (lead 1)	0.74	1.90	1.161	0.16	2697
#Purchased apps (lead 2)	0.43	0.54	0.106	0.42	2697
#Purchased apps (lead 3)	0.19	0.74	0.543	0.33	2697
#Purchased apps (lead 4)	0.18	0.26	0.072	0.22	2697
#Purchased apps (lead 5)	0.24	0.32	0.086	0.49	2697
#Purchased apps (lead 6)	0.17	0.30	0.130	0.12	2697
#Purchased apps (lead 7)	0.15	0.21	0.056	0.37	2697
#Purchased apps (lead 8)	0.24	0.32	0.084	0.22	2697
#Newborn apps (lead 1)	0.41	0.75	0.347	**0.02**	2697
#Newborn apps (lead 2)	0.32	0.68	0.359	**0.01**	2697
#Newborn apps (lead 3)	0.37	0.64	0.266	**0.02**	2697
#Newborn apps (lead 4)	0.34	0.68	0.339	**0.03**	2697
#Newborn apps (lead 5)	0.44	0.65	0.209	0.12	2697
#Newborn apps (lead 6)	0.40	0.68	0.272	**0.03**	2697
#Newborn apps (lead 7)	0.51	0.77	0.260	0.15	2697
#Newborn apps (lead 8)	0.40	0.69	0.292	0.08	2697
#App updates (lead 1)	1.31	1.84	0.532	0.06	2697
#App updates (lead 2)	1.01	2.22	1.217	0.10	2697
#App updates (lead 3)	1.14	1.80	0.663	**0.04**	2697
#App updates (lead 4)	1.28	1.62	0.340	0.32	2697
#App updates (lead 5)	1.24	3.84	2.596	0.17	2697
#App updates (lead 6)	1.01	1.87	0.858	**0.01**	2697
#App updates (lead 7)	1.19	3.54	2.358	0.18	2697
#App updates (lead 8)	1.24	3.55	2.316	0.16	2697
#Removed app (lead 1)	0.32	1.12	0.798	**0.00**	2697
#Removed app (lead 2)	0.19	0.43	0.236	**0.01**	2697
#Removed app (lead 3)	0.21	0.37	0.164	**0.01**	2697
#Removed app (lead 4)	0.15	0.29	0.142	**0.00**	2697
#Removed app (lead 5)	0.17	0.35	0.179	**0.01**	2697
#Removed app (lead 6)	0.19	0.30	0.112	**0.04**	2697
#Removed app (lead 7)	0.16	0.30	0.133	**0.02**	2697
#Removed app (lead 8)	0.13	0.26	0.132	**0.02**	2697
#Dismissed apps (lead 1)	0.21	0.49	0.276	**0.01**	2697
#Dismissed apps (lead 2)	0.16	0.29	0.124	**0.06**	2697
#Dismissed apps (lead 3)	0.18	0.31	0.131	**0.03**	2697
#Dismissed apps (lead 4)	0.13	0.23	0.100	**0.00**	2697
#Dismissed apps (lead 5)	0.13	0.27	0.141	**0.00**	2697
#Dismissed apps (lead 6)	0.15	0.26	0.108	**0.04**	2697
#Dismissed apps (lead 7)	0.14	0.25	0.115	**0.02**	2697
#Dismissed apps (lead 8)	0.11	0.21	0.106	**0.04**	2697
#Sold apps (lead 1)	0.11	0.63	0.522	**0.00**	2697
#Sold apps (lead 2)	0.03	0.14	0.112	**0.02**	2697
#Sold apps (lead 3)	0.03	0.07	0.033	0.08	2697
#Sold apps (lead 4)	0.02	0.06	0.042	**0.01**	2697
#Sold apps (lead 5)	0.04	0.07	0.038	0.37	2697
#Sold apps (lead 6)	0.03	0.04	0.005	0.79	2697
#Sold apps (lead 7)	0.03	0.05	0.018	0.49	2697
#Sold apps (lead 8)	0.02	0.05	0.026	0.12	2697

Notes: This table presents the balance check results for the treatment and control groups using a 1:1 matching ratio with Mahalanobis distance for the 2-period lag before the M&A events. The comparison covers various app-related metrics to assess whether the matching achieves balance across the key variables. The reported mean differences and p-values indicate whether there are statistically significant differences between the matched groups after controlling for observable characteristics. Only the 2-period lag balance check is shown here, while the complete balance checks for additional periods are provided in the subsection 3.A.6. Variables with p-values greater than or equal to 0.05 are highlighted in bold. Please refer to Table 3.1 for detailed descriptions of the variables.

Table 3.17: Balance check after matching -continue

PSM(1:1)					
	Control	Treatment	Mean-diff	P-value	N
Controls					
Log(#Owned apps) (lead 1)	1.90	1.96	0.059	0.16	2697
Log(#Owned apps) (lead 2)	1.82	1.87	0.044	0.31	2697
Log(#Owned apps) (lead 3)	1.76	1.81	0.041	0.35	2697
Log(#Owned apps) (lead 4)	1.72	1.75	0.035	0.43	2697
Log(#Owned apps) (lead 5)	1.68	1.70	0.022	0.62	2697
Log(#Owned apps) (lead 6)	1.63	1.65	0.017	0.71	2697
Log(#Owned apps) (lead 7)	1.58	1.59	0.012	0.79	2697
Log(#Owned apps) (lead 8)	1.53	1.52	-0.007	0.87	2697
Log(1+#Ratings) (lead 1)	6.57	6.80	0.233	**0.04**	2697
Log(1+#Ratings) (lead 2)	6.46	6.69	0.226	0.05	2697
Log(1+#Ratings) (lead 3)	6.38	6.60	0.216	0.06	2697
Log(1+#Ratings) (lead 4)	6.30	6.49	0.192	0.10	2697
Log(1+#Ratings) (lead 5)	6.22	6.39	0.174	0.13	2697
Log(1+#Ratings) (lead 6)	6.13	6.30	0.163	0.16	2697
Log(1+#Ratings) (lead 7)	6.02	6.17	0.157	0.18	2697
Log(1+#Ratings) (lead 8)	5.91	6.03	0.121	0.30	2697
Log(1+#Paid apps) (lead 1)	0.17	0.22	0.052	**0.02**	2697
Log(1+#Paid apps) (lead 2)	0.16	0.21	0.047	**0.03**	2697
Log(1+#Paid apps) (lead 3)	0.16	0.20	0.039	0.07	2697
Log(1+#Paid apps) (lead 4)	0.16	0.20	0.036	0.10	2697
Log(1+#Paid apps) (lead 5)	0.16	0.19	0.032	0.14	2697
Log(1+#Paid apps) (lead 6)	0.16	0.19	0.030	0.16	2697
Log(1+#Paid apps) (lead 7)	0.16	0.19	0.032	0.13	2697
Log(1+#Paid apps) (lead 8)	0.16	0.19	0.033	0.12	2697
Log(1+#Translated des) (lead 1)	0.74	0.77	0.031	0.42	2697
Log(1+#Translated des) (lead 2)	0.72	0.75	0.027	0.49	2697
Log(1+#Translated des) (lead 3)	0.70	0.73	0.028	0.46	2697
Log(1+#Translated des) (lead 4)	0.68	0.71	0.033	0.38	2697
Log(1+#Translated des) (lead 5)	0.67	0.70	0.026	0.49	2697
Log(1+#Translated des) (lead 6)	0.65	0.68	0.029	0.42	2697
Log(1+#Translated des) (lead 7)	0.63	0.66	0.023	0.52	2697
Log(1+#Translated des) (lead 8)	0.62	0.63	0.008	0.81	2697
Log(1+#Miss downloads) (lead 1)	0.05	0.10	0.052	**0.00**	2697
Log(1+#Miss downloads) (lead 2)	0.05	0.06	0.015	0.14	2697
Log(1+#Miss downloads) (lead 3)	0.05	0.07	0.017	0.11	2697
Log(1+#Miss downloads) (lead 4)	0.05	0.07	0.015	0.19	2697
Log(1+#Miss downloads) (lead 5)	0.05	0.06	0.014	0.18	2697
Log(1+#Miss downloads) (lead 6)	0.04	0.06	0.018	0.09	2697
Log(1+#Miss downloads) (lead 7)	0.05	0.06	0.014	0.21	2697
Log(1+#Miss downloads) (lead 8)	0.05	0.06	0.015	0.18	2697
Average rating (lead 1)	3.60	3.64	0.039	0.26	2697
Average rating (lead 2)	3.61	3.67	0.056	0.10	2697
Average rating (lead 3)	3.61	3.67	0.056	0.11	2697
Average rating (lead 4)	3.63	3.68	0.051	0.15	2697
Average rating (lead 5)	3.64	3.68	0.037	0.30	2697
Average rating (lead 6)	3.66	3.69	0.031	0.39	2697
Average rating (lead 7)	3.66	3.68	0.019	0.62	2697
Average rating (lead 8)	3.65	3.67	0.021	0.59	2697
#Status updates (lead 1)	1.82	3.18	1.352	**0.00**	2697
#Status updates (lead 2)	1.50	3.16	1.665	0.05	2697
#Status updates (lead 3)	1.65	2.70	1.047	**0.01**	2697
#Status updates (lead 4)	1.74	2.55	0.814	0.05	2697
#Status updates (lead 5)	1.90	4.75	2.851	0.15	2697
#Status updates (lead 6)	1.57	2.80	1.232	**0.00**	2697
#Status updates (lead 7)	1.81	4.53	2.721	0.14	2697
#Status updates (lead 8)	2.25	4.51	2.253	0.22	2697

causal estimates by controlling for observable characteristics. Full balance checks for additional periods are included in this section to provide a detailed assessment of the matching process.

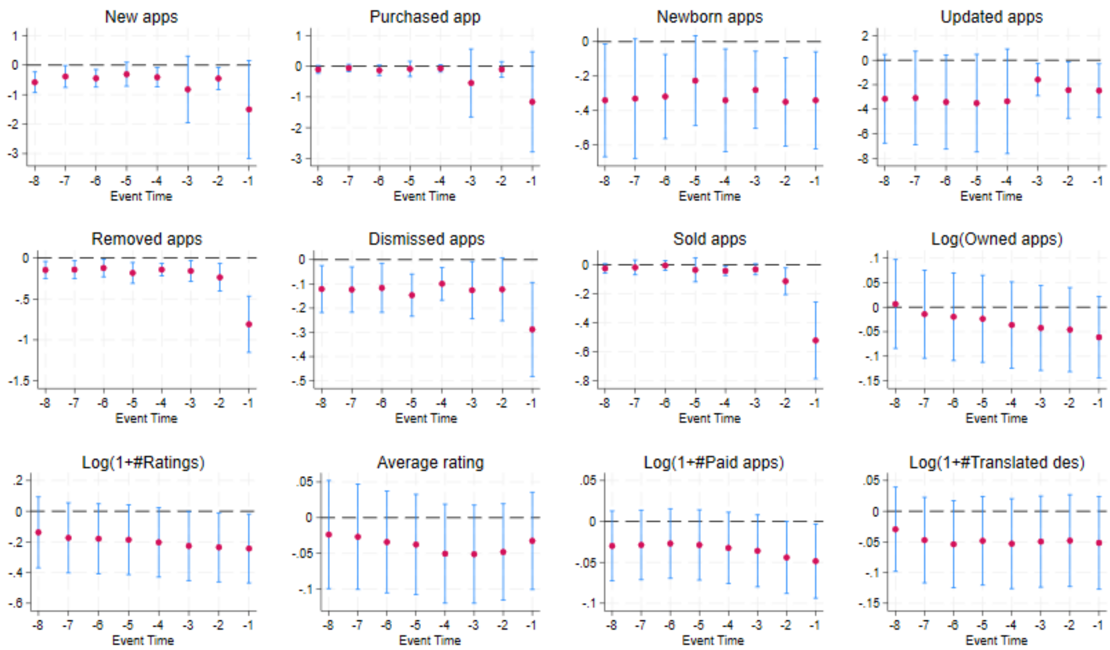

Figure 3.9: Balance check

Notes: The figure displays balance checks for multiple treatment outcomes across event time periods preceding merger and acquisition (M&A) events. The x-axis represents event time, with negative values indicating pre-treatment periods. The variables examined include new apps, newborn apps, app updates, removed apps, dismissed apps, sold apps, the log of owned apps, the log of total ratings, average rating, the log of non-free apps, and the log of apps with translated descriptions. The balance checks pertain to a specification that considers 8 pre-treatment periods. For greater detail and the balance checks for the other periods, please refer to the subsection 3.A.5.

Figure 3.9 presents balance checks over various pre-treatment periods corresponding to Table 3.16 and Table 3.17, spanning eight time periods prior to the M&A event. For most outcomes, including new apps, purchased apps, app updates, and the log of total ratings, the coefficients remain close to zero across the pre-treatment periods, with overlapping confidence intervals, indicating no significant differences between the treated and control groups before the event. A few outcomes, such as removed apps, dismissed apps, and sold apps, show minor deviations from zero in certain periods. Overall, the balance check suggests there is no systematic divergence between the treated and control groups during the pre-treatment periods, confirming the absence of confounding trends that could bias the analysis of post-treatment effects. This supports the validity of the matching procedure and ensures comparability between groups.

3.A.7 Placebo test

In order to ensure the robustness of our findings, we conducted an in-time placebo test. This test shifts the treatment time forward by several periods to simulate a 'fake' treatment event, and examines whether significant effects are observed in periods when no real treatment occurred.

Using the following process: We create multiple fake treatment periods by moving the actual treatment time forward by 1 to 8 periods. For each shifted period, we estimate the treatment effect using the Callaway and Sant'Anna (2021) DiD estimator. We apply this to key app development outcome variables: new app creation, app updates, app removals, purchased apps, dismissed apps, and monetization changes. We then plot the placebo effect estimates and their confidence intervals, expecting these estimates to be statistically insignificant if the original treatment effects are robust.

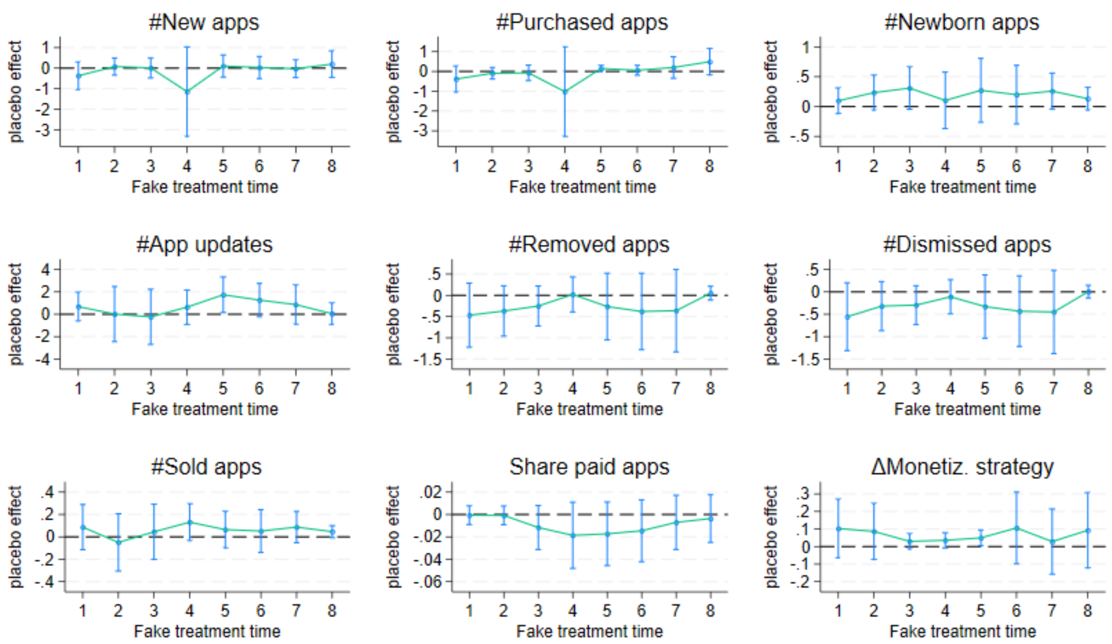

Figure 3.10: Placebo test

3.A.8 Main result using quarterly Japanese and US data

Table 3.18: Summary of M&A by event date of Japanese Google play market

No.	Event Date	Sum MA	Count Merger	Count Acquisition
1	2016-04-04	154	16	138
2	2016-07-04	146	21	125
3	2016-10-03	130	19	111
4	2017-01-02	157	16	141
5	2017-04-03	178	26	152
6	2017-07-03	170	34	136
7	2017-10-02	188	18	170
8	2018-01-01	199	29	170
	Sum	1322	179	1143

Note: The event date represents the M&A that occurred in the preceding six week interval. For example, the event date '2016-04-04' refers to M&A that took place between January 2016 and April 2016. In this context, a merger refers to the combination of two app developers to form a new app developer, while an acquisition occurs when one app developer acquires another, resulting in the target developer no longer existing as an independent entity. The earliest M&A activity within the observation window is used to define the event date.

To verify the robustness of our findings, we extend our analysis to the Japanese app market using quarterly data from Google Play. The M&A events identified follow the same methodology as in our main US six-week sample, providing a cross-national perspective on M&A effects in the app market.

Of the M&A events identified in the Japanese sample, 80.4% correspond with those observed in the US sample, suggesting a significant overlap and reinforcing the impact of major acquisitions across different markets. Table 3.18 presents a summary of M&A activity by event date. Table 3.19 provides the summary statistics for Japnaniese Google Play data.

To validate the robustness of our main findings from the US data, we replicate the analysis with Japanese market data, as shown in Table 3.20. The Japanese results align with the US findings, showing similar coefficient signs across core outcome variables, including reduced new app entries, purchased apps, and Share paid apps, even where statistical significance is lower. This consistency suggests that M&A events lead to app consolidation and shifts in monetization models across markets. However, due to the

Table 3.19: M&A vs Non-M&A summary statistics 2015-19

	(M&A)		(Non M&A)			(Mean diff)
	Mean	SD	Mean	SD	Coeff.	(t-value)
Outcome variables						
(1)#New apps	3.53	44.33	0.29	5.50	3.25^{a}	(9.14)
(1a)#Purchased apps	1.00	9.27	0.06	2.27	0.93^{a}	(12.59)
(1b)#Newborn apps	2.54	43.25	0.56	6.23	1.98^{a}	(5.71)
(2)#App updates	4.06	35.98	0.39	5.01	3.67^{a}	(12.73)
(3)#Removed apps	2.28	29.57	0.19	4.63	2.09^{a}	(8.81)
(3a)#Dismiss apps	1.89	28.27	0.19	4.59	1.70^{a}	(7.52)
(3b)#Sold apps	0.39	7.96	0.01	0.53	0.38^{a}	(6.01)
(4)Share paid apps	0.01	0.08	0.06	0.19	-0.04^{a}	(-65.30)
(5)ΔMonetization strategy	0.07	1.09	0.02	0.85	0.05^{a}	(6.31)
Control variables						
Log(#Owned apps)	1.86	1.28	0.63	0.91	1.23^{a}	(119.80)
Log(1+#Ratings)	6.75	3.17	3.27	2.45	3.48^{a}	(136.87)
Log(1+#Paid apps)	0.20	0.57	0.10	0.36	0.10^{a}	(21.04)
Log(1+#Translated des)	1.82	1.17	1.00	0.71	0.81^{a}	(86.71)
Log(1+#Miss downloads)	0.06	0.32	0.02	0.17	0.04^{a}	(17.15)
Average rating	3.70	0.90	3.51	1.51	0.19^{a}	(26.55)
#Status updates	6.73	57.13	0.93	7.83	5.80^{a}	(12.67)
N observations	8,712,040		15,587			8,727,627

Notes: The superscript ^{a} denotes statistical significance at the 1% level. This table compares summary statistics for developers involved in mergers or acquisitions (M&A) with those who were not (Non-M&A) using Japan Quarterly Google play data. Row definitions refer to Table 3.1

Table 3.20: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT (JAPAN)	-0.310	**-0.880*****	0.832	0.963	-0.152	-0.0883	-0.0640	**-0.007*****	-0.210
	(1.031)	(0.209)	(0.997)	(1.385)	(0.699)	(0.649)	(0.197)	(0.00239)	(0.136)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	25046	25046	25046	25046	25046	25046	25046	25046	25046
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES
ATT (US)	-0.100	-0.144	0.177	**1.017***	0.101	-0.0214	**0.123****	**-0.007*****	0.016
	(0.476)	(0.464)	(0.111)	(0.572)	(0.141)	(0.120)	(0.060)	(0.003)	(0.036)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	78500	78500	78500	78500	78500	78500	78500	78500	78500
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: The table reports ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 4 pre-treatment periods, where the matching is applied to units with at least 4 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

quarterly nature of the Japanese data (as opposed to the six-week intervals in the US data), there are fewer pre-event periods for matching, which limits statistical power and may reduce the significance of certain outcomes. This cross-market comparison nonetheless underscores the generalizability of the non-price effects of M&A in the app market.

3.A.9 Main result using various segment matching

To further ensure the robustness of our findings, we conducted an additional matching procedure that assigns treated units to control units from different market segments. Initially, we categorized app developers based on the size of their primary category within Google Play's approximately 49 app categories (For further details, please refer to subsection 3.3.5). This process allowed us to identify each developer's dominant market based on their largest share of market activity. In the main analysis, our matching approach involved pairing treated units (developers involved in mergers or acquisitions) with control units that operated within the same market segment, thereby ensuring comparability within a shared market context.

For this robustness check, we modify our matching criteria to pair treated units with control units that do not belong to the same dominant market segment. This approach tests whether the observed effects of M&A are consistent even when comparing across distinct market segments, providing insight into whether our main findings are driven by market-specific factors or broader strategic behaviours. By matching treated developers with controls from different segments, we aim to evaluate whether the impact of M&A on app-related outcomes holds when cross-market differences are considered, thus reinforcing the validity and generalizability of our results.

Table 3.21: Event study ATT results

	(1) #New apps	(1a) #Purchased apps	(1b) #Newborn apps	(2) #App updates	(3) #Removed apps	(3a) #Dismissed apps	(3b) #Sold apps	(4) Share paid apps	(5) ΔMonetiz. strategy
ATT	**-1.024*****	**-0.819*****	**-0.166***	1.039	**-0.882*****	**-0.579*****	**-0.303*****	**-0.005***	0.006
	(0.272)	(0.250)	(0.098)	(0.680)	(0.192)	(0.175)	(0.095)	(0.002)	(0.013)
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Developer FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	51277	51277	51277	51277	51277	51277	51277	51277	51277
PT†	YES	YES	YES	YES	YES	YES	YES	YES	YES

Notes: The table reports ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8 pre-treatment periods, where the matching is applied to units with at least 8 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results from Table 3.21, which use cross-market matching, confirm the robustness of our main findings. Specifically, we continue to find significant reductions in new app launches (-1.926), purchased apps (-1.762), and removed apps (-0.968), alongside significant declines in dismissed (-0.681) and sold apps (-0.287). At the same time, there is a significant increase in app updates (1.186), as well as smaller but still significant effects on changes in monetization strategy (0.025) and the share of paid apps (-0.004). These effects mirror the patterns observed in our primary within-market analysis (see Table 3.6), indicating that the impact of M&A on app development, portfolio management, and monetization persists even when treatment and control groups are matched across different market segments. Although some effect sizes differ slightly in magnitude, the consistency of statistical significance across key outcomes highlights that our core conclusions are not driven by within-market dynamics alone. This further underscores the generalizability of our findings and demonstrates that M&A-induced strategic shifts among app developers extend across diverse market contexts.

3.A.10 Heterogeneous effect of app updates

In app market, not all app updates are created equal: some represent substantial innovations or strategic repositioning, while others are routine maintenance. Accurately distinguishing between these types of updates is essential for analysing the heterogeneity of post-M&A innovation effects. To operationalize this, we classify app updates into 'Major' and 'Minor' updates based on the content of their release notes (`whats_new`). This section details our methodological approach, which leverages state-of-the-art large language models for scalable and consistent classification.

3.A.11 Classification pipeline

We employ a two-stage approach combining NLP and manual rule design, implemented as follows:

Data preparation: We extract the free-text update logs for all observed app updates. All logs are pre-processed to ensure consistent encoding, and non-English entries are automatically translated into English prior to classification.

Automated classification with GPT: To systematically categorize each update, we utilize a state-of-the-art large language model (Azure OpenAI, GPT-4o) to assign each update log to one of eight predefined categories: *first release*, *new feature*, *major improvement*, *bug fix*, *performance improvement*, *UI/UX change*, *security update*, and *other*. For each batch of logs, the model returns a category for every update, following a strict, mutually exclusive coding scheme.

Mapping to 'Major' versus 'Minor' Updates: For empirical analysis, we collapse these granular categories into two broad classes: **Major updates:** updates classified as *first release*, *new feature*, or *major improvement*. These reflect substantial innovation or strategic product shifts. **Minor updates:** updates classified as *bug fix*, *performance improvement*, *UI/UX change*, *security update*, or *other*. These typically represent incremental or maintenance-oriented changes.

Each update is assigned a binary indicator for 'key innovation' accordingly.

3.A.12 Implementation details

The classification is conducted using the Azure OpenAI API, with the following procedure: Update logs are processed in batches for efficiency. The GPT model is prompted with explicit category definitions to ensure consistency and reduce ambiguity. For non-English logs, automatic translation precedes classification, guaranteeing all updates are judged on comparable semantic grounds.

This methodology offers several key advantages: **Scalability:** Capable of classifying millions of update logs in a manner infeasible for manual coding. **Consistency:** Application of a fixed prompt and classification scheme mitigates coder subjectivity and time-based drift. **Content sensitivity:** By utilizing full-text logs, the classification captures semantic nuance (e.g., distinguishing a major redesign from a routine bug fix) that would be missed by keyword-based or metadata approaches.

3.A.13 Empirical analysis - app developer level

The resulting indicators for key and Minor updates are integrated into our event study framework, enabling the estimation of heterogeneous treatment effects of M&A on distinct types of innovation activity. As shown in Table 3.22, we report separate average treatment effects for key and Minor updates, revealing nuanced impacts of M&A on app-level innovation behaviour.

Table 3.22: Event study ATT results

	(1) #App updates	(1a) #Major updates	(1b) #Minor updates
ATT	**1.145****	0.355	**0.790*****
	(0.514)	(0.302)	(0.275)
Time FE	✓	✓	✓
Developer FE	✓	✓	✓
Observations	73876	73876	73876
PT†	YES	YES	YES

Notes: The table reports ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8 pre-treatment periods, where the matching is applied to units with at least 8 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.22 presents the estimated ATT of M&A on overall app updates, Major updates, and Minor updates. The results indicate that M&A are associated with a statistically significant increase in the total number of app updates (ATT = 1.180), confirming that post-M&A periods see a notable rise in developer activity aimed at refining or improving their app portfolios. Disaggregating by update type, the increase is predominantly driven by a significant rise in Minor updates (ATT = 0.824), while the effect on Major updates is positive but not statistically significant (ATT = 0.356). This suggests that, following an M&A event, developers are more likely to engage in routine maintenance or incremental improvements, rather than introducing major new features or overhauls. The lack of a significant effect on key innovations implies that M&A may prioritize portfolio stability and continuous refinement over disruptive product innovation in the immediate aftermath. These findings highlight the importance of distinguishing between different types of updates when evaluating the innovative consequences of consolidation in digital markets.

3.A.14 Empirical analysis - app level

Table 3.23: Event study ATT results

	(1) #App updates	(1a) #Major updates	(1b) #Minor updates
ATT	**0.095*****	**0.038*****	**0.057*****
	(0.003)	(0.002)	(0.003)
Time FE	✓	✓	✓
Developer FE	✓	✓	✓
Observations	546373	546373	546373
PT†	YES	YES	YES

Notes: The table reports ATT results from applying the Callaway and Sant'Anna (2021) estimator to the data, examining the impact of M&A on various app-related outcomes. The standard errors are clustered by the app developer ID. The chi-squared statistic and its p-value are reported in parentheses. The models include time fixed effects (Time FE) and developer fixed effects (Developer FE). PT† represents whether the parallel trend assumption holds. The results are shown for 8 pre-treatment periods, where the matching is applied to units with at least 8 pre-treatment periods, respectively. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.23 reports the estimated ATT of M&A on app updates, Major updates, and Minor updates at the app level. The findings indicate that, following an M&A event, the likelihood of an app being updated increases by 0.095 (significant at the 1% level). Importantly, both key and Minor updates exhibit statistically significant increases, with ATT of 0.038 and 0.057, respectively, each significant at the 1

These results provide clear evidence that M&A spur innovation-related activity not only at the developer portfolio level but also for individual apps. The positive and significant effects for both key and Minor updates suggest that, post-M&A, acquired or merged apps are more likely to undergo maintenance, incremental improvements, and-though to a lesser degree-major updates or feature introductions. While the magnitude of the effect is somewhat larger for Minor updates, the increase in Major updates is also robustly significant, contrasting with the developer-level analysis where only Minor updates showed a significant effect.

This nuanced pattern suggests that the innovation response to M&A is not confined to routine maintenance activity but also involves more intensive development effort at the app level. In particular, the app-level results indicate that, following consolidation, developers reallocate resources toward both minor and major updates, reflecting changes in development priorities within existing portfolios. The observed increase in major updates may be associated with post-merger integration, feature consolidation, or alignment with new strategic objectives, rather than with expansion through new app creation.

Overall, these results underscore the value of disaggregating innovation measures and provide novel evidence that M&A in digital markets can stimulate both incremental and major product improvements for individual apps. This reinforces the view that post-M&A innovation is multi-dimensional, encompassing both routine and transformative changes across the app portfolio.

Chapter 4

Text-defined markets: An NLP framework for antitrust with evidence from the app economy[1]

We study two questions central to competition analysis in digital markets: how to delineate relevant markets when user prices are often zero, and how mergers within those markets affect non-price conduct. Using Google Play data from 2015–2019, we construct text-defined markets from app descriptions and update notes, measure concentration, and identify horizontal overlap between merging developers. We estimate merger effects using a matched, staggered DiD design. We identify 1,914 horizontal mergers with overlapping presence in the same text-defined market, around one-fifth of which are flagged by EU or US concentration screens. Flagged horizontal mergers are associated with substantial portfolio reallocation, including higher app exits (about 6 percentage points) and modest increases in update activity, with no detectable price effects. Horizontal mergers below policy thresholds show smaller adjustments. By contrast, non-horizontal mergers are associated with higher update rates (around 5 percentage points), modest price increases (about \$0.12), and selective portfolio pruning. Overall, consolidation in zero-price digital markets

[1]This chapter is independent work by the author. The data were provided by Michael Kummer. I developed the NLP methodology, constructed the text-defined markets, implemented the empirical analysis, and wrote the chapter. Franco Mariuzzo and Michael Kummer provided comments and suggestions at various stages. All remaining errors are my own. This manuscript benefited from the use of AI-assisted writing tools for grammar and style improvements; all intellectual content, analysis, and interpretations remain the sole responsibility of the author. We thank participants at the internal seminar at the School of Economics (UEA) for helpful comments.

primarily reshapes portfolio composition and maintenance effort, while price effects arise mainly in non-horizontal integrations.

4.1 Introduction

The relationship between market structure and innovation has long been central in industrial organization (Arrow, 1972; Schumpeter J A, 1942). Recent work emphasizes how these dynamics unfold in digital, innovation-intensive settings where diffusion, intangibles, and platform design shape productivity and competitive pressure (Kalyani et al., 2025; Brynjolfsson et al., 2019). In parallel, new evidence highlights how acquisitions can eliminate nascent competition (Cunningham et al., 2020) and how leaders' strategic control over idea commercialization shapes observed innovation (Argente et al., 2020).

This paper contributes to the literature by studying non-price competitive outcomes around M&A in the mobile-app economy and by developing an empirically grounded, text-based approach to market definition when traditional price signals are weak or absent.

Our research question is twofold. First, how should 'relevant markets' be delineated in settings with a high prevalence of zero posted prices on at least one side, multi-sided interactions, and rapid product turnover? Second, how do M&A affect non-price outcomes including entry, product updates (incremental innovation/maintenance), and portfolio consolidation in digital markets? Classic merger analysis focuses on prices, elasticities, and diversion ratios (Farrell and Shapiro, 1990; Werden and Froeb, 1994; Berry et al., 1995). Yet in digital markets, quality, privacy, and innovation frequently mediate competitive pressure, and the standard SSNIP logic can be ill-posed on a single side of the platform (Rochet and Tirole, 2003; Filistrucchi et al., 2014). We therefore pair a method-first contribution-constructing 'text-defined markets' from product descriptions and change logs using modern NLP-with an application that evaluates strategic responses around M&A in the app economy.

To operationalize this approach, we construct relevant markets using a transparent, multi-stage NLP pipeline designed to balance economic interpretability and robustness. App descriptions and update notes are embedded into a semantic vector space, reduced in dimension, and clustered into candidate markets using two complementary algorithms: HDBSCAN, which flexibly captures uneven competitive structure, and a partitioning method (K-means), which imposes more regular market boundaries. To assess sensitivity to market granularity, we implement both algorithms under alternative caps on cluster size, ranging from 300 to 1,200 apps per market. We evaluate the stability of resulting market definitions using standard overlap metrics-the ARI and Jaccard similarity-ensuring that empirical results do not hinge on a particular clustering choice. This design yields text-defined markets that are economically interpretable, scalable, and robust to reasonable variation in methodological assumptions.

Empirically, we (i) embed app descriptions and update notes, (ii) reduce dimensionality, and (iii) cluster products into coherent 'text-defined markets', validating clusters and computing concentration (e.g., HHI). We then leverage modern DiD estimators for staggered timing, combining propensity-score based matching with event-study aggregation (Callaway and Sant'Anna, 2021). Our evaluation aligns with contemporary merger frameworks that emphasize transparent, design-based identification and interpretable competitive metrics (Miller and Weinberg, 2017). For policy interpretation, we adopt US structural-presumption thresholds (post-merger HHI $\geq$ 1,800 and ΔHHI $\geq$ 100; with a separate share-based screen) and define EU 'screened-in' markets as those outside the EU safe harbours.

We report two main findings. First, text-defined markets formed from app descriptions and update notes enable an economically meaningful analysis of competition: among 1,914 horizontal mergers, roughly one-fifth are flagged by EU or US concentration screens (most by both), while about 90% of all transactions are non-horizontal. Second, these classifications allow to pin down increases in market-power with sharply different post-merger conduct. Flagged horizontal transactions-those exceeding standard HHI and ΔHHI thresholds-are associated with pronounced portfolio rationalization: next-period exits rise by about 6 percentage points ('pruning') and update activity increases modestly, with no detectable price change. Non-flagged horizontals prune less intensely (about 4 percentage points) and show limited innovation response, whereas non-horizontal M&A events do not affect exit, but are characterized by higher update rates (about 5 percentage points) and small positive price adjustments (about \$0.12). These patterns, estimated in staggered DiD and event-study designs with app and time fixed effects, suggest consolidation in already concentrated markets reallocates effort within portfolios rather than raising prices, while cross-market acquisitions result in price increases and foster incremental innovation. The results are robust across clustering methods and policy regimes.

Conceptually, we provide a portable, evidence-first market-definition framework which is particularly useful for zero or negligible-price, multi-sided environments: text-defined markets operationalize substitution without relying on prices or static industry codes, yet enable the quantification of market concentration using widely-used measures (e.g., HHI) and dynamic market-level event studies. Substantively, we bring new evidence on non-price competitive effects of M&A in a large digital ecosystem, distinguishing horizontal from non-horizontal M&A to document competing effects in entry versus incremental innovation. Methodologically, we integrate modern NLP with state-of-the-art DiD estimators to produce policy-relevant, design-based estimates where structural demand or price data are unavailable.

The structure of the paper is as follows: Section 4.2 reviews the relevant literature. Section 4.3 details the text-based market-construction pipeline; Section 4.4 introduces the data and concentration measures; Section 4.5 presents the DiD and event-study results; Section 4.6 concludes. Appendices report clustering diagnostics, ARI/Jaccard stability, and recovery of policy flags under alternative market definitions.

4.2 Literature review

4.2.1 Market definition: foundations and digital-market challenges

The delineation of relevant markets is a cornerstone of competition analysis, and a fundamental ingredient for the assessment of market power and merger effects. Classical approaches, codified in the US DOJ/FTC Horizontal Merger Guidelines and the European Commission's Market Definition Notice, rest on the principle that a market comprises the smallest set of products and geographies over which a hypothetical monopolist could impose a SSNIP profitably. This logic has roots in early industrial organization work: Kenneth G. Elzinga and Thomas F. Hogarty (1973) trade-flow tests formalised geographic boundaries via 'little in from outside' and 'little out from inside' criteria, while Stigler et al. (1985) proposed using price co-movements to infer integration. The development of demand-estimation techniques, notably Berry et al. (1995), enabled empirical implementation of the hypothetical monopolist test by simulating the profitability of coordinated price increases over candidate product sets. Subsequent refinements, such as Werden (1997) integration of critical-elasticity analysis and diversion-based measures, and Katz and Shapiro (2002) linkage of critical loss to aggregate diversion ratios, grounded market definition in measurable elasticities and substitution patterns, while clarifying potential pitfalls such as the reverse Cellophane fallacy (Froeb and Werden (1992)).

This classical canon agrees on two key points. First, market definition is an empirical exercise in identifying the set of close substitutes that constrain pricing, which can be informed by a variety of evidence including elasticities, diversion, product characteristics, and customer perceptions. Second, while concentration metrics like the HHI are often reported, they are meaningful only when the underlying market boundaries reflect actual competitive constraints.

Digital markets-where posted prices on the user side are often zero or negligible - are often characterized by multi-sided interactions, and strong network effects-complicate this framework. In such settings, SSNIP tests on a single side may be ill-posed: for example, in zero-price services such as search engines or social media, users ‘pay’ in attention or data rather than money. Scholars have proposed adapting the test to quality or non-price dimensions, such as the Small but Significant Non-transitory Decrease in Quality (SSNDQ)(Newman (2015)), or applying a ‘net-price’ test across both sides of a platform (Rochet and Tirole (2003) ; Filistrucchi et al. (2014)). Caillaud and Jullien (2003) and Armstrong (2006) highlight that multi-homing and single-homing patterns critically affect substitution: platforms serving predominantly single-homing users may constitute distinct markets even when switching costs are low in principle, whereas high multi-homing can integrate platforms into a common market.

Indirect network effects imply that demand on one side of a multi-sided platform depends on participation on the other, so changes in participation or quality on one side can alter competitive constraints on the other (Condorelli and Padilla, 2020). In such environments, competition is often mediated through non-price dimensions—such as product variety, quality, and innovation-rather than posted prices alone. These features complicate the application of standard price-based market-definition tools and motivate approaches that infer substitution directly from product characteristics rather than observed prices.

In response, recent guidance from the Commission (2024) emphasises that, in digital markets, relevant market definition should account for non-price competition, the characteristics of multi-sided platforms (including cross-side effects), and dynamic competitive processes. The Competition and Markets Authority (2021) also recognises that non-price factors can be critical, and-in some cases-market definition may need to reflect the presence of multi-sided platforms. Agencies increasingly complement formal market-definition exercises with direct evidence of competitive effects—such as diversion ratios or upward pricing pressure—especially where rigid boundaries risk obscuring the mechanisms of harm (Kaplow (2010); Farrell and Shapiro (2010)). For empirical work, this evolution underscores the importance of flexible, evidence-rich approaches that remain anchored in the economic principle of substitution while accommodating the multi-dimensional, dynamic nature of competition in the digital economy.

4.2.2 Empirical approaches in practice: evidence and gaps

Empirical competition analysis has long relied on quantitative tools for defining markets, measuring concentration, and assessing the competitive effects of mergers. Classical contributions established a coherent framework grounded in demand substitution, market concentration metrics, and the modelling of post-merger incentives. Landes and Posner (1997) provided a formal link between market share, demand elasticity, and market power, offering a theoretical rationale for using concentration as an initial screening tool. Hall and Tideman (1967) formalised properties of concentration measures (including HHI), which later became standard in merger guidelines. The development of structural demand estimation, notably Berry et al. (1995), enabled empirical implementation of the hypothetical monopolist test by recovering substitution patterns from product characteristics and market data, thus providing the empirical basis for merger simulation (Werden and Froeb (1994)). These methods were operationalised in case studies such as Nevo (2001), which demonstrated how scanner data and structural demand models could delineate narrow product markets and simulate competitive effects.

Classical merger analyses emphasised unilateral pricing incentives and the role of efficiencies. Farrell and Shapiro (1990) formalised the conditions under which mergers without cost savings raise prices, laying the foundation for upward pricing pressure indices. Williamson (1972) illustrated the potential for efficiencies to offset anti-competitive harm, a theme empirically tested in merger retrospectives such as Ashenfelter and Hosken (2010), which found that most studied mergers resulted in higher consumer prices, and Blonigen and Pierce (2016), which detected post-merger increases in markups without corresponding productivity gains. Together, these studies reinforced that empirical practice must move beyond concentration thresholds to incorporate direct evidence on substitution patterns, price effects, and efficiencies.

Recent developments in digital markets challenge core assumptions of the classical toolkit. The economics of multi-sided platforms (Rochet and Tirole (2003)) and the recognition of cross-group network effects (Filistrucchi et al. (2014)) are consistent with competitive constraints cannot be assessed by analysing one side of the market in isolation. zero or negligible-price markets further complicate empirical application of the SSNIP test, prompting proposals for quality-based analogues (Evans (2011)) and highlighting the relevance of non-price competition dimensions such as privacy and innovation (Gal and Rubinfeld (2016)). In such environments, conventional concentration measures may be misleading: network effects can accelerate market tipping, while common ownership across rivals (Azar et al. (2018)) can reduce competitive intensity without changing market structure in the traditional sense.

Methodological innovations have expanded empirical tools for market delineation. Hoberg and Phillips (2016) demonstrated that NLP of firm disclosures can identify product-market boundaries that differ from static industrial classifications, offering a dynamic and data-driven view of competition. Extending this approach, Gugler et al. (2024) used sentence embeddings and clustering to define 'text-based markets' in the app economy, allowing concentration and competitive dynamics to be measured where price data are absent. These methods align with the classical principle of defining markets through substitution, but replace the reliance on price and category codes with high-dimensional textual similarity, enabling application to rapidly evolving digital sectors.

Despite these advances, empirical practice exhibits notable gaps. First, the integration of non-price dimensions into market definition remains underdeveloped in applied work, even where policy guidance recognises their importance. Second, potential competition effects, such as 'killer acquisitions' of nascent rivals (Cunningham et al. (2020)), remain difficult to quantify ex ante within standard empirical frameworks. Third, while NLP and other machine-learning methods offer scalable market definitions, their adoption in enforcement practice is still at a formative stage, with questions about transparency, interpretability, and robustness. Addressing these gaps will require combining the empirical rigour of structural industry organisation with the flexibility of modern data-driven approaches, ensuring that evidence remains aligned with the economic principle of substitution while accommodating the complexities of digital competition.

4.3 Text embedding and clustering methodology

This study examines the impact of M&A on the non-price outcomes of app developers, with a particular focus on innovation and product differentiation as reflected in app descriptions. To analyse latent patterns in these descriptions and track their evolution around M&A events, we employ a multi-stage NLP and clustering pipeline tailored to the scale and multilingual nature of the Google Play dataset.

4.3.1 Text preprocessing and language identification

All app descriptions are subjected to rigorous preprocessing, including the removal of emojis, special characters, URLs, emails, and numbers. This ensures a consistent textual format, minimizing noise and enhancing comparability across developers and over time. Language detection is performed using the `langdetect` library, enabling us to assign each description to the appropriate embedding model and to account for the multilingual composition of the app market.

4.3.2 Semantic embedding of app descriptions

We represent the semantic content of each app description using the `SentenceTransformer` library.[2] For English texts, we employ the `all-MiniLM-L6-v2` model,[3] while non-English descriptions are encoded with the `paraphrase-multilingual-MiniLM-L12-v2` model.[4] Each description is mapped into a dense vector in a high-dimensional semantic space, capturing relationships in meaning beyond simple keyword overlap. To reduce dimensionality while preserving neighbourhood structure, we apply Uniform Manifold Approximation and Projection (UMAP) with parameters tuned to dataset size within each language–category subset.

4.3.3 Clustering algorithms and design choices

Market construction proceeds by clustering the UMAP-reduced embeddings. We implement two recursive algorithms that reflect complementary approaches to uncovering latent market structure: HDBSCAN): A density-based algorithm that flexibly identifies clusters of varying shapes and densities. For each language-category block, HDBSCAN is applied to UMAP embeddings. Clusters exceeding a size threshold are recursively re-embedded and split until either the cluster size falls below the threshold or a maximum depth is reached. Noise points are reassigned to their nearest clusters, ensuring near-complete coverage. We conduct parallel analyses with thresholds of 300, 600, 900, and 1200 apps per cluster to assess robustness to this design choice. Recursive k-means clustering algorithm (k-means): A partitioning algorithm that imposes spherical boundaries. We apply K-means to each UMAP-reduced block, with the number of clusters chosen adaptively. Oversized clusters are recursively split until the imposed size threshold is satisfied. We conduct parallel analyses with thresholds of 300, 600, 900, and 1200 apps per cluster to assess

[2]The `SentenceTransformer` framework implements transformer-based language models that map sentences or short documents into fixed-length vector representations optimised for semantic similarity tasks, such as clustering and nearest-neighbour search.

[3]`all-MiniLM-L6-v2` is a lightweight sentence-embedding model trained using knowledge distillation from larger transformer architectures. It is optimised to capture semantic similarity while remaining computationally efficient, making it suitable for large-scale text embedding tasks.

[4]`paraphrase-multilingual-MiniLM-L12-v2` is a multilingual sentence-embedding model trained on parallel and paraphrase data across more than 50 languages, enabling comparable semantic representations across languages.

robustness to this design choice.

We deliberately restrict attention to HDBSCAN and K-means. Alternative methods such as Density-Based Spatial Clustering of Applications with Noise (DBSCAN) or agglomerative clustering were not pursued, as they scale poorly to millions of observations or generate unstable results in high-dimensional text embeddings. Our choice balances economic interpretability, computational feasibility, and robustness.

4.3.4 Clustering pipeline and robustness

Clustering is performed separately for each language–category subset, producing clusters defined by the combination of language, category, and algorithm–threshold specification. This design reflects the heterogeneity of the app economy and prevents spurious mixing of unrelated products. To evaluate robustness, we compare outcomes across clustering methods and thresholds, using overlap metrics such as the ARI and Jaccard similarity, as well as distributions of cluster sizes. This ensemble approach confirms that our main findings are not artifacts of a single algorithm or parameter choice, but reflect stable underlying patterns in the data.

Quantitative stability diagnostics-ARI and Jaccard similarity across algorithms and size caps-are reported in Appendix 4.A.1. These metrics show high consistency, particularly for the HDBSCAN (900-cap) specification, which therefore serves as our baseline market definition in the analyses that follow. Further construction artifacts—including cluster size distributions across caps (Figure 4.6) and extended baseline properties (Table 4.8)-are reported in Appendix 4.A.2.

4.4 Data

4.4.1 Clustering the app space and baseline specifications

Markets and units of analysis. We define a market m as a text-defined cluster (category $\times$ language) at date t arising from our NLP+clustering pipeline (HDBSCAN, k-Nearest Neighbours (k-NN)). Within each market–time cell, we compute concentration (i) across apps and (ii) across developers.

Table 4.1: Sample coverage and language composition

	Apps	
	Count	**Share (%)**
Total apps	7,888,829	100.0
English apps	7,315,093	92.7
Non-English apps	573,736	7.3
Categories	49	–

Notes: The table reports the language composition and overall coverage of the processed Google Play dataset (2015–2019). Each observation corresponds to a unique app identifier with at least one valid textual description and developer ID. Language classification is based on automatic detection using the `langdetect` library; English accounts for 92.7% of all descriptions. 'Categories' refers to the 49 official Google Play categories represented in the data.

Our empirical analysis relies on a novel panel dataset of mobile applications released on Google Play between 2015 and 2019. For each app, we observe developer identity, assigned category, language, release dates, update logs, and textual descriptions. Because descriptions evolve over time, we standardize on a medium-period snapshot of each app's description. This provides a stable and representative textual record for clustering while retaining the panel structure of the data for the applied part of the study, where we examine dynamics such as entry, product updates, and withdrawals.

Table 4.1 summarizes overall sample coverage and language composition. The processed dataset contains nearly 7.9 million distinct apps spanning 49 unique categories. Language is observed for all apps, with English accounting for 92.7 percent of the sample. The remaining 7.3 percent (over half a million apps) are in other languages, highlighting the multilingual nature of the platform.

While English dominates overall, there is substantial heterogeneity across categories. Some categories—such as Family, Photography, and Racing-are almost exclusively English, while others-notably Books & Reference, Music & Audio, and

Comics-feature a higher prevalence of non-English apps. These differences reflect category-specific demand and cultural orientation, and motivate clustering separately by category and language to avoid conflating products that target distinct user bases.

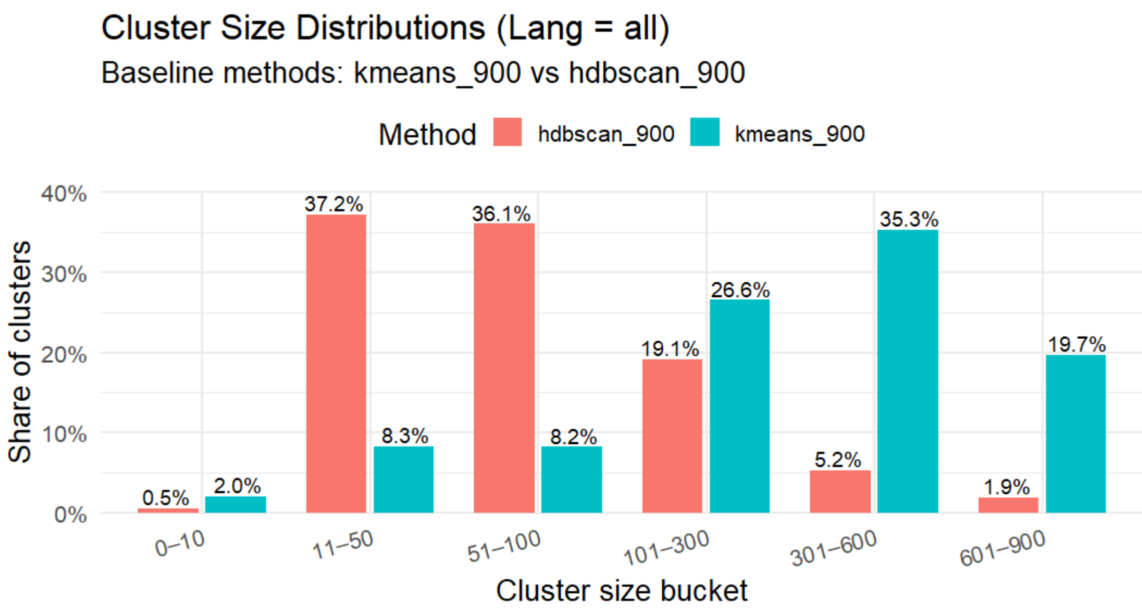

Figure 4.1: Cluster size distributions for baseline text-defined markets

Notes: The figure reports the distribution of cluster sizes for baseline markets constructed with HDBSCAN (`hdbscan_900`) and k-means (`kmeans_900`), pooling all languages. The horizontal axis shows cluster-size buckets (number of apps per cluster), and the vertical axis shows the share of clusters in each bucket. HDBSCAN yields many small clusters (11–100 apps), while k-means produces fewer very small clusters and a substantially larger mass of medium and large clusters (301–900 apps).

Figure 4.1 illustrates how HDBSCAN and k-means produce different cluster-size distributions: HDBSCAN yields many small, dense markets, whereas k-means partitions the space into fewer but substantially larger markets. While the two algorithms generate different size distributions, the key market properties relevant for subsequent analysis—such as concentration levels, the identification of horizontal overlaps, and the share of markets flagged under policy screens-remain qualitatively stable across methods and cluster-size caps. Appendix 4.A.3 reports detailed robustness checks using both HDBSCAN and k-means across caps (300–1200), confirming that our main empirical results do not hinge on a particular clustering specification.

4.4.2 Construction of market measures and policy screens

App-level concentration. For each app i in market m and time t, let w_{imt} denote its cumulative rating stock—the running total of ratings the app has accumulated up to t. Because Google Play does not report installs or usage, we use this stock as a proxy for the app's *relative demand intensity* within the market, consistent with prior work in settings where direct install or usage data are unavailable: an app that has attracted more ratings has, all else equal, been used and downloaded more. We turn stocks into market shares by normalizing each app's stock by the total stock in its market,

$$s_{imt} = \frac{w_{imt}}{\sum_j w_{jmt}},$$

so that shares sum to one within each market and capture each app's share of cumulative user engagement. These shares are the inputs to the concentration measures below; the rating stock therefore enters the HHI only through the shares s_{imt}.

Market concentration is then measured using the HHI and four-firm concentration ratio (CR4):

$$\mathrm{HHI}_{mt} = 10{,}000 \sum_i s_{imt}^2, \qquad \mathrm{CR4}_{mt} = 100 \sum_{i=1}^{4} s_{(i)mt},$$

where $s_{(i)mt}$ denotes the i-th largest market share.

Developer-level concentration. We next measure concentration across developers within each market m and time t. Ownership is tracked dynamically, assigning each app to its current developer based on the merger event ledger (the acquirer replaces the pre-event owner from the event date g onward). Let f index developers and i apps, and denote by $\mathcal{J}_{fmt}$ the set of apps owned by developer f in market m at time t. The developer's cumulative stock is

$$\mathrm{stock}_{fmt} = \sum_{i \in \mathcal{J}_{fmt}} \mathrm{stock}_{imt}, \qquad s_{fmt}^{\mathrm{dev}} = \frac{\mathrm{stock}_{fmt}}{\sum_{f'} \mathrm{stock}_{f'mt}},$$

where stock_{imt} is the monotone cumulative rating stock of app i.

Developer-level concentration is summarized by the HHI and CR4:

$$\mathrm{HHI}^{\mathrm{dev}}_{mt} = 10{,}000 \sum_{f} \left(s^{\mathrm{dev}}_{fmt}\right)^2, \qquad \mathrm{CR4}^{\mathrm{dev}}_{mt} = 100 \sum_{k=1}^{4} s^{\mathrm{dev}}_{(k)mt},$$

where $s^{\mathrm{dev}}_{(1)mt} \geq s^{\mathrm{dev}}_{(2)mt} \geq \ldots$ denote ordered developer shares. For completeness, we also record the number of active apps $N_{\mathrm{apps},mt}$, developers $N_{\mathrm{dev},mt}$, and total market stock $W_{mt} = \sum_f \mathrm{stock}_{fmt}$.

Policy screens. To benchmark concentration changes against established antitrust criteria, we apply the screening thresholds used by the main competition authorities in the US and the EU. These rules, summarized in Table 4.2, identify markets that exceed standard post-merger concentration or change-in-concentration cutoffs.

Under the 2023 US *Merger Guidelines* (U.S. Department of Justice and Federal Trade Commission, 2023), a structural presumption arises when the post-merger HHI exceeds 1,800 and increases by at least 100 points, with an additional share-based screen when the merged share exceeds 30%. The 2004 EU *Horizontal Merger Guidelines* (European Commission, 2004) define 'safe harbours' under which mergers are unlikely to raise concerns; markets outside these bounds are labelled 'EU screened-in'.

Table 4.2: Policy screen summary

Jurisdiction	**Screen (flags)**
US (2023)	Structural: $\mathrm{HHI}^{\mathrm{post}} \geq 1{,}800$ & $\Delta\mathrm{HHI} \geq 100$ Share-based: merged share $\geq 30\%$ & $\Delta\mathrm{HHI} \geq 100$
EU (2004)	Safe harbours: $\mathrm{HHI}^{\mathrm{post}} < 1{,}000$; or $1{,}000 \leq \mathrm{HHI}^{\mathrm{post}} < 2{,}000$ & $\Delta\mathrm{HHI} < 250$; or $\mathrm{HHI}^{\mathrm{post}} \geq 2{,}000$ & $\Delta\mathrm{HHI} < 150$. Screened-in = none apply.

4.4.3 Market structure in text-defined app level

Table 4.3: Market concentration and average cluster size

Statistic	**All**	**English**	**Non-English**
Number of markets (clusters)	74,512	68,024	6,488
Median HHI (0–10,000)	4,556	4,512	5,140
HHI interquartile range [p25, p75]	[2,777, 7,446]	[2,749, 7,379]	[3,493, 7,686]
90th percentile of HHI	9,603	9,571	9,651
Median CR4 (pp)	95.3	95	97.7
Average cluster size (apps per market)	106	107	85

Notes: Markets are text-defined HDBSCAN clusters. Concentration is computed from *stock* shares based on cumulative ratings within each market × date. For each market we take the *time median* of HHI and $CR4$, then report the cross-market distribution. Apps with zero stock mass at (m, t) are excluded from share calculations. The *Average cluster size* is a static property of the clustering: the mean number of unique apps assigned to each market, and is independent of time.

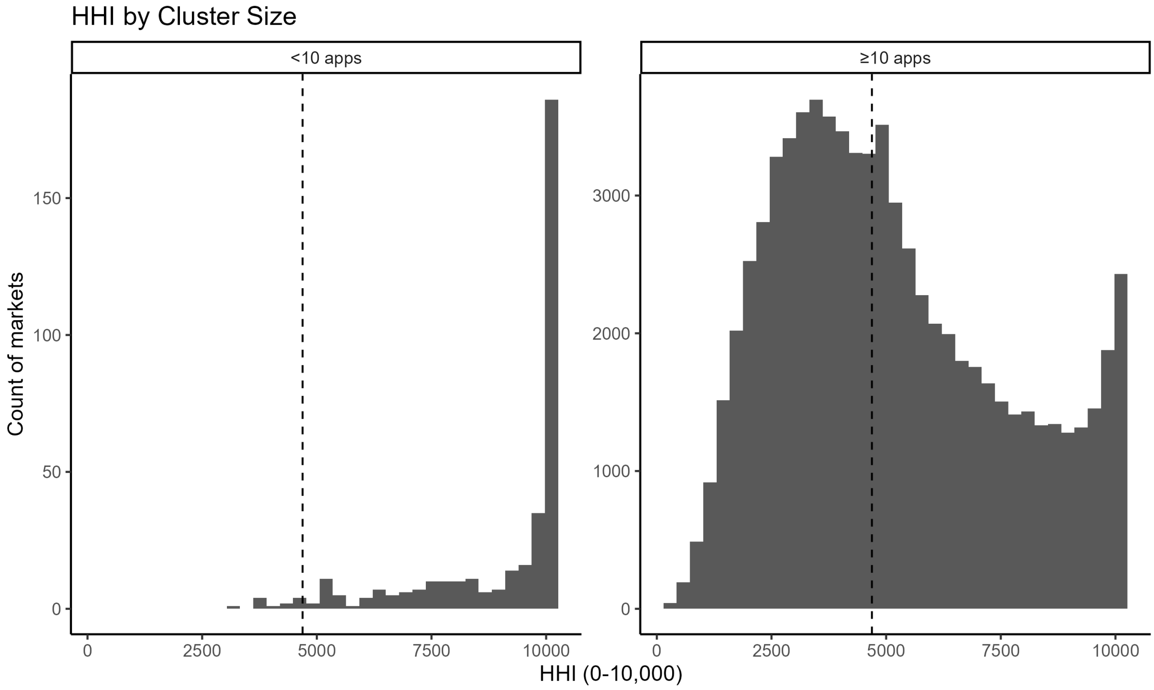

Figure 4.2: Cluster-level concentration (app-level), split by cluster size

Notes: Each panel shows the distribution of HHI for text-defined markets built via HDBSCAN (max cluster size cap = 900 apps). For every cluster, we compute the HHI at the app level using cumulative ratings *stocks* and then take the median across time. Panels split clusters by size: '$<$ 10 apps' and '$\geq$ 10 apps,' where size is the number of distinct apps mapped to the cluster in the taxonomy. The dashed vertical line marks the overall median standardized HHI across all clusters in the sample. In the $<$ 10 apps panel, the spike near 100% reflects very small markets that are frequently dominated by a single app. In the $\geq$ 10 apps panel, the hump around 50% reflects *effective duopolies*-two leading apps account for most of the cumulative stock even when additional apps are present with small shares.

Table 4.3 reveals a strikingly concentrated structure in the app economy. Across all text-defined markets, the median HHI is 4,556 on the 0–10,000 scale-implying that, in a typical market, the equivalent of roughly two firms account for the entire cumulative user base. Concentration is only slightly lower among English-language markets (median = 4,512) and notably higher among non-English markets (median = 5,140), suggesting that smaller or more linguistically segmented markets tend to be dominated by fewer players. The interquartile range spans from about 2,800 to 7,400, and the 90th percentile exceeds 9,600, indicating that at least 10 percent of markets approach monopoly conditions. Consistent with these high HHI, the median

CR4 reaches 95 percent, meaning that four leading apps capture virtually the entire cumulative stock of ratings in the median market.

The shape of the distribution in Figure 4.2 underscores how concentration varies systematically with market size. Clusters containing fewer than ten apps exhibit a pronounced mass near 100 percent, reflecting near-monopolistic dominance within small, specialized niches. Even among larger markets with ten or more apps, the density peaks around 50 percent, consistent with effective duopolies where two major apps divide most of user engagement while peripheral competitors command negligible shares. Economically, these patterns point to strong winner-take-most dynamics typical of digital ecosystems: network effects, user lock-in, and platform-mediated visibility allow a handful of developers to capture disproportionate market share, leaving limited scope for fringe entrants despite the vast number of available apps.

4.4.4 Market structure in text-defined developer level

Table 4.4: Market concentration

Statistic	**All**
Number of markets (clusters)	74,726
Median HHI (0–10,000)	4,091
HHI interquartile range [p25, p75]	[2,387, 6,918]
90th percentile of HHI	9,643
Median developers per market	12
Developers per market IQR [p25, p75]	[5, 25]

Notes: Markets are fixed HDBSCAN text clusters (static app to cluster mapping). HHI statistics are computed per market×date cell and summarized with *equal weight across cells*. 'Developers per market' counts distinct developers with positive stock in the cell.

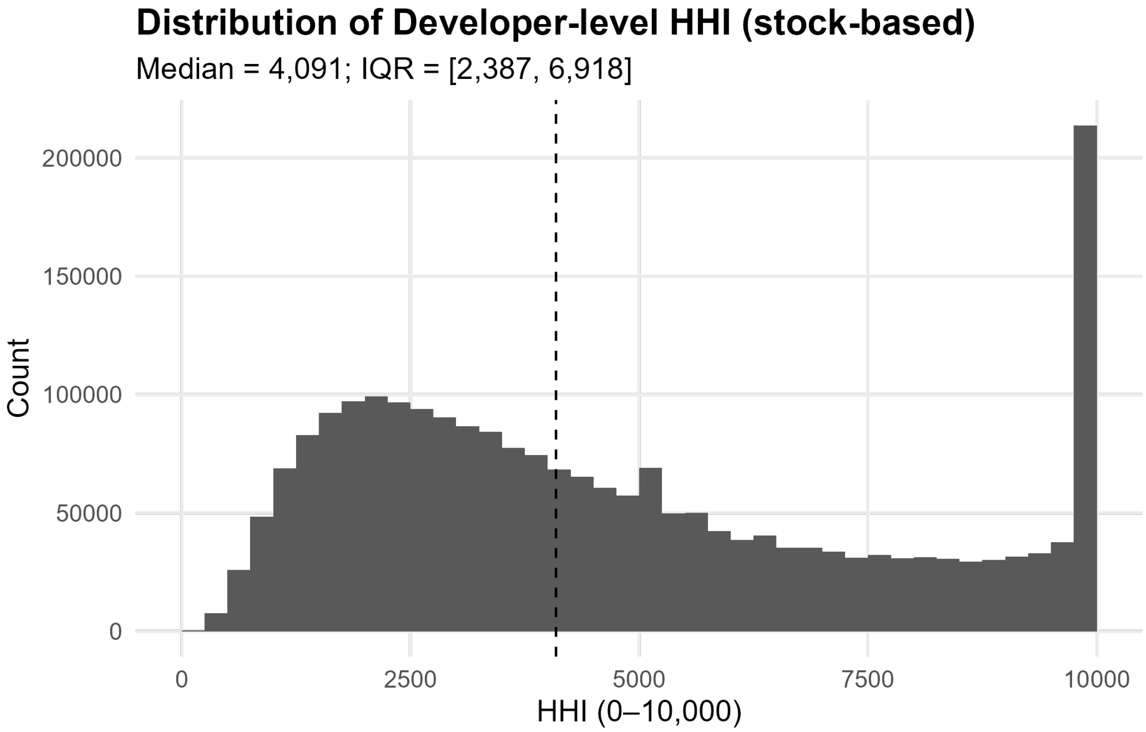

Figure 4.3: Distribution of the event-time change in developer-level concentration (ΔHHI_L)

Notes: The figure reports the distribution of ΔHHI_L-the change in developer-level concentration at the event time t_0 relative to the 6-period pre-event average-computed for every event×cluster pair (each deal mapped to all HDBSCAN markets it touches).

The patterns in Table 4.4 and Figure 4.3 are consistent with two distinct mechanisms. First, when merging developers both operate within the same text-defined market, developer-level concentration rises mechanically as overlapping product portfolios are combined under common ownership, directly increasing the acquirer's market share. Second, beyond this immediate ownership effect, post-merger portfolio rationalization- through coordinated updates, maintenance, or promotional emphasis-provides a channel through which user attention and installed base may be reallocated toward the merged developer over time, even in the absence of immediate entry or exit. To make these mechanisms operational, the subsequent analysis defines markets as fixed HDBSCAN text clusters with a static app→cluster mapping and computes developer-level shares from stock mass (cumulative rating counts) aggregated across all apps owned by the developer in a market. An *overlap* (horizontal exposure) is recorded at the developer level for an event–market pair whenever the survivor and at least one counterparty both have strictly positive pre–event stock $W_{jmt} > 0$ in that same market at t_0^-; deals are mapped to every market they touch, and ΔHHI_L is computed per event×market as in Figure 4.3.

Interpreted through this lens, the level facts—median HHI^{dev} of 4,091 on a 0–10,000 scale with a median of 12 active developers per market in Table 4.4—imply concentrated but not monopolistic competition, in which multi–title portfolios diffuse single–app dominance at the firm level. The distribution of the event-time concentration change ΔHHI_L in Figure 4.3 is centered near zero with a fat right tail, indicating that most M&A occur without developer-level overlap in the relevant market (hence small average shifts), while the subset of overlapping cases delivers economically meaningful increases in concentration precisely where combined portfolios expand effective reach. This definition of overlap will anchor our horizontal–versus–non-horizontal split in the results that follow, ensuring that comparisons are tied to pre–t_0 market presence rather than post-merger outcomes.

4.4.5 Concentration measures and policy flags

Table 4.5: HHI-based policy screens under two market definitions: text-defined markets and a benchmark category market

	Text-defined markets (HDBSCAN, $cap = 900$)		Benchmark: Category market (49 Categories)	
Panel A. Non-horizontal M&A (no overlapping market)	18,727	90.7	520	63.0
Panel B. Horizontal M&A - policy screening within overlapping markets				
EU competition policy				
Flagged	333	17.4	0	0.0
Not flagged	1,581	82.6	305	100.0
Total (EU horizontal)	1,914	100.0	305	100.0
US (DOJ/FTC structural-presumption test)				
Flagged	359	18.8	0	0.0
Not flagged	1,555	81.2	305	100.0
Total (US horizontal)	1,914	100.0	305	100.0
Panel C. Overlap composition among screened-in				
EU flagged (only)	16	4.3	–	–
US flagged (only)	42	11.2	–	–
Both (EU ∩ US)	317	84.5	–	–
Total flagged (EU ∪ US)	375	100.0	–	–

Notes: The left block uses text-defined markets from HDBSCAN clustering ($cap = 900$) at t_0^-, with developer shares built from monotone rating stocks; Panel B applies EU and US thresholds within the overlapping (horizontal) market. The right block is a benchmark based on category markets (dominant Google Play category at t_0^-) from Chapter 3. Because this coarse benchmark aggregates heterogeneous rivals, it poorly isolates true overlaps at event time; as a result, nearly all horizontal cases appear 'Not flagged,' and flagged-overlap composition (Panel C) is not informative (shown as '–').

Table 4.5 shows that, under text-defined markets, the overwhelming majority of deals are non-horizontal: 90.7% (18,727 developer–events) occur with no overlap in the relevant market, compared with only 63.0% (520) when using the coarse category benchmark. Among the remaining horizontal cases identified by the text-based definition, 17.4% are screened in by the EU thresholds and 18.8% by the US structural presumption, with nearly identical complements not flagged (82.6% and 81.2%, respectively). By contrast, the category benchmark flags 0% of horizontal cases under either regime, implying a 17–19 percentage point gap in the detected rate of potentially concerning overlaps. The intersection among screened-in cases is large-84.5% of flagged overlaps satisfy both EU and US criteria-indicating substantial cross-jurisdiction concordance once overlap is measured in economically meaningful product space.

These contrasts are consistent with two mechanisms central to the paper's market-definition approach. First, text-defined markets segment competition at the level of substitutable functionalities and usage language, avoiding the aggregation bias implicit in broad categories; this precision both raises the share of deals correctly classified as non-horizontal (by excluding unrelated rivals) and concentrates the truly horizontal mass in markets where shares and ΔHHI are policy-relevant. Second, even within overlap, the fact that roughly four-fifths of event–market pairs are not flagged suggests that many mergers combine portfolios with diffuse pre-merger positions or generate modest immediate ΔHHI-consistent with multi-title scope and intra-developer substitution that temper mechanical concentration jumps. Together, these facts imply that text-defined markets sharpen the screen by ruling out spurious 'overlaps' and isolating the 17–19% of horizontal cases where concentration concerns are most likely to arise, while the category benchmark systematically understates potential competitive interactions by diluting head-to-head rivalry across heterogeneous apps.

4.4.6 Policy context and threshold sensitivity

Tzanaki (2023) argues that the EU's high, turnover-based jurisdictional screens created a 'killer acquisitions' enforcement gap, prompting a de facto lowering of jurisdictional gates via a recalibrated use of Article 22 EU Merger Regulation (EUMR) and, in parallel, added transparency through the Digital Markets Act (DMA)'s Article 14 reporting-before the Court's *Illumina/Grail* judgment narrowed that path and pushed more weight back to national 'call-in' powers and transaction-value tests. In spirit, that institutional debate mirrors our *screening* exercise: the share of horizontal overlaps flagged is mechanically increasing in the tightness of the cut-offs. In our text-defined markets (HDBSCAN, *cap*=900), 17.4% (EU) and 18.8% (US) of horizontal cases are flagged at current HHI/ΔHHI rules (Table 4.5); modest relaxations-e.g., lowering the US ΔHHI trigger from 100 to 80 (holding $HHI \geq 1800$) or treating HHI>1800 *or* ΔHHI>200 as screened-in on the EU side-would pull a non-trivial mass of near-threshold overlaps into the flagged set. Substantively, this is consistent with our finding that most deals are non-horizontal, while the horizontal subset features concentrated yet diffuse pre-merger positions that generate modest ΔHHI on average, leaving meaningful density just below today's policy lines.

4.4.7 Descriptive statistics: treated apps versus matched controls

Table 4.6 summarizes the key outcome variables for treated and control apps, highlighting systematic pre-event differences that are economically meaningful in this setting. Treated apps exhibit substantially higher update intensity: the mean probability of any timely update is 0.169 for treated apps versus 0.078 for controls, an absolute gap of 0.091. The composition of that activity also skews toward more consequential releases: 'Major update' updates average 0.047 for treated apps against 0.014 for controls, while 'Minor update' are 0.122 versus 0.064. By contrast, listed prices are lower among treated apps (USD 0.174) than among controls (USD 0.364), a difference of USD 0.190, or about 52% of the control mean. Churn-related outcomes are similar at baseline: the next-period dismissal rate is identical (0.034 in both groups), and transfers ('Sold following period') are rare in both samples, with treated apps slightly less likely to be sold (0.003 versus 0.004).

These patterns are consistent with acquirers targeting more actively maintained and innovation-intensive products, even if those apps monetize less through posted prices—plausibly reflecting freemium or advertising-based models typical in consumer mobile markets. The absence of baseline differences in dismissal rates suggests comparable short-run attrition risk across groups, while the marginally lower pre-event sale rate among treated apps is consistent with tighter pre-merger

Table 4.6: Summary statistics of app–level outcomes for treated and control samples

Outcome	$N_{\mathbf{obs}}$	$N_{\mathbf{ids}}$	**Mean**	**SD**	**Max**
Treated apps					
(1) Price ($)	388,260	19,126	0.174	1.787	135.75
(2) App update	388,260	19,126	0.169	0.375	1
(2a) Major update	388,260	19,126	0.047	0.211	1
(2b) Minor update	388,260	19,126	0.122	0.327	1
(3) Dismissed following period	382,284	19,126	0.034	0.182	1
(4) Sold following period	369,134	19,126	0.003	0.050	1
Control apps					
(1) Price ($)	337,946	19,564	0.364	2.777	200
(2) App update	337,946	19,564	0.078	0.268	1
(2a) Major update	337,946	19,564	0.014	0.117	1
(2b) Minor update	337,946	19,564	0.064	0.245	1
(3) Dismissed following period	331,417	19,395	0.034	0.182	1
(4) Sold following period	320,045	19,051	0.004	0.063	1

Notes: N_{obs} is the number of app×date observations with non-missing outcomes; N_{ids} is the number of distinct apps with at least one non-missing observation. Treated apps are those directly exposed to a merger event at the app level; control apps ($G = 0$) are drawn from the matched developer pool established in Chapter 3, restricted to developers that are never treated and excluding any app that ever appears in a treated set. Outcome variables are defined as follows: *Price*-listed price in US dollars. The vast majority of apps in the sample are free, so price effects are identified from a relatively small subset of paid apps and should be interpreted accordingly. ; *App update*-indicator for any version update within the time; *Major update*-indicator for major feature or design changes inferred from update notes; *Minor update*-indicator for minor maintenance or bug-fix updates; *Dismissed following period*-indicator equal to 1 if the app becomes unavailable in the following time; and *Sold following period*-indicator equal to 1 if the app is transferred to a different developer. Both 'Dismissed following period' and 'Sold following period' are rare binary outcomes. As such, treatment effects estimated in subsequent difference-in-differences analyses should be interpreted as changes in the incidence of these events rather than literal percentage-point changes in probability.

ownership or lower propensity to divest. Because control apps in Table 4.6 are drawn from the matched developer pool established in Chapter 3 and our empirical designs include time and app fixed effects, these level differences primarily inform interpretation rather than identification: they highlight that treatment is concentrated in higher-maintenance, lower-price product lines, a selection margin that helps contextualize the post-merger dynamics estimated in the event-study analyses.

4.4.8 Robustness to market-definition choices

Robustness to market-definition choices. To assess whether our findings hinge on the specific clustering configuration, we recompute policy screens under alternative market definitions. Using HDBSCAN (cap = 900) as the baseline, other HDBSCAN parametrizations recover roughly 60–75% of the baseline policy flags, indicating stability to parameter tuning. Switching to K-means at the same granularity still recovers about half of the EU and US flagged apps, suggesting that the main policy classifications are broadly robust to the choice of algorithm (see Table 4.9 in Appendix 4.A.3).

4.5 Empirical results

We construct the empirical sample in two stages. First, treated units are defined at the app level by mapping each M&A into its relevant text-defined markets at the event time and applying the concentration screens described above. Transactions are classified as Non-horizontal M&A when the merging developers have no overlapping market presence, and as Horizontal M&A when overlap exists in at least one text-defined market. Within the horizontal group, we distinguish between Flagged and Not flagged cases based on standard policy thresholds: flagged if HHI $\geq$ 1,800 and ΔHHI $\geq$ 100 under the US structural-presumption rule, or outside the EU safe harbours (i.e., screened-in) under EU guidelines. This produces disjoint sets of treated apps corresponding to Non-horizontal, Horizontal – Flagged, and Horizontal – Not flagged transactions. In the second stage, we build a panel for each policy group comprising (i) treated apps exposed at t_0 (retaining the first exposure when multiple events occur) and (ii) control apps drawn from the matched developer sample established in Chapter 3, restricted to developers that are never treated and excluding any app that ever appears in a treated set. Event time is defined at the app level, and we track six outcomes—price, overall update activity, Major updates, Minor updates, dismissals, and sales—to estimate staggered DiD effects using the Callaway and Sant'Anna (2021) estimator, allowing for unbalanced panels, clustering standard errors at the app level, and reporting both dynamic paths and aggregated effects.

Table 4.7: ATT by US policy screen

	US Flagged	US Not Flagged	Non-horizontal
App outcomes			
(1)Price ($)	0.002 (0.029)	0.089*** (0.026)	0.117*** (0.018)
(2)App update	0.023** (0.012)	-0.005 (0.008)	0.047*** (0.003)
(2a)Major update	0.008 (0.005)	-0.012** (0.005)	0.014*** (0.002)
(2b)Minor update	0.015 (0.011)	0.007 (0.007)	0.032*** (0.003)
(3)Dismissed following period	0.063*** (0.009)	0.041*** (0.002)	0.035*** (0.001)
(4)Sold following period	-0.624*** (0.018)	-0.494*** (0.012)	-0.222*** (0.003)
Time FE	Yes	Yes	Yes
App FE	Yes	Yes	Yes
Apps (treated), N_{tr}	731	1,863	16,544
Apps (controls), N_{ctrl}	19,564	19,564	19,564
Observations (N_{obs})	353,169	372,440	676,667

Notes: Entries report overall ATT from the staggered DiD estimator of Callaway and Sant'Anna (2021), computed separately for three US policy groups defined at the event time: *Flagged* (structural presumption met: HHI $\geq$ 1,800 and ΔHHI $\geq$ 100), *Not Flagged* (at least one of the two thresholds not met), and *Non-horizontal*. Each cell shows the point estimate with clustered standard errors in parentheses; standard errors are clustered at the app level. All specifications include time and app fixed effects. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Treated/control app counts by group are taken from the overall (price) specification: US Flagged $N_{\text{tr}} = 731$, Not flagged $N_{\text{tr}} = 1{,}863$, Non-horizontal $N_{\text{tr}} = 16{,}544$; controls $N_{\text{ctrl}} = 19{,}564$. Outcome-specific samples vary slightly due to missing-ness: e.g., for 'Dismissed' and 'Sold,' the controls are 19,395 and 19,051.

The estimates in Table 4.7 quantify within-developer changes relative to matched controls, with time and app fixed effects absorbed. As shown in Figure 4.4, pre-treatment coefficients are generally small, but not uniformly zero across all policy-defined groups, which is unsurprising given the staggered timing and narrow treated cells. We therefore view the post-treatment responses as indicative of merger exposure, while acknowledging that identification is weaker for those groups exhibiting pre-event movement. The resulting pattern still supports a re-optimization story that is sharper for horizontal, policy-flagged mergers than for non-horizontal integrations.

Flagged horizontal mergers where concentration screens bind exhibit a clear portfolio-rationalization pattern. The incidence of any app update increases by 0.023, indicating intensified maintenance activity on retained titles. At the same time, exit and divestiture margins adjust sharply: the incidence of next-period dismissals rises by 0.063, while app sales collapse, with an estimated ATT of -0.624, reflecting a near elimination of post-merger divestitures. Benchmarked against baseline means in Table 4.6, the dismissal effect corresponds to an increase in exit incidence of roughly 185 percent relative to the pre-merger mean of 0.034, while the update response amounts to about 14 percent of the baseline update rate (0.169). Prices, by contrast, display no economically meaningful change. Taken together, these patterns point to active internal reallocation within merged portfolios greater maintenance of core products combined with selective pruning of peripheral apps-rather than post-merger monetization through higher posted prices.

By contrast, non-flagged horizontal mergers exhibit little adjustment along innovation margins. Overall update activity remains essentially unchanged (-0.005), with only a modest reduction in Major updates (-0.012). Nevertheless, these transactions still involve active portfolio streamlining: the incidence of next-period dismissals increases by 0.041, and app sales decline sharply (ATT = -0.494), indicating a substantial contraction in post-merger divestiture activity even when concentration screens do not bind.

The most expansive responses arise for non-horizontal mergers. Update activity increases by 0.047, with both Major and Minor updates rising (0.014 and 0.032, respectively), alongside a pronounced increase in prices of 0.117. Dismissals also rise by 0.035, pointing to selective pruning within expanding portfolios. Relative to baseline means, the increase in updates amounts to roughly 28 percent of the pre-merger update rate, while the price response corresponds to about 67 percent of the average listed price. Together, these patterns are consistent with scope-driven expansion and monetization synergies when developers combine complementary, non-overlapping product lines, rather than with competitive responses driven by increased horizontal market power.

Appendix 4.A.4 reports the corresponding estimates under the EU policy screens together with extended dynamic specifications. The results closely mirror the US patterns: concentration-screened mergers exhibit intensive post-merger rationalization and higher exit rates, whereas non-horizontal integrations display sustained increases in innovation and moderate price growth. These parallel findings reinforce the interpretation that the observed differences are structural rather than sample-specific.

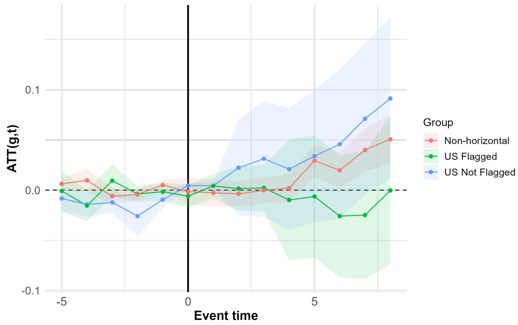

(a) Price ($)

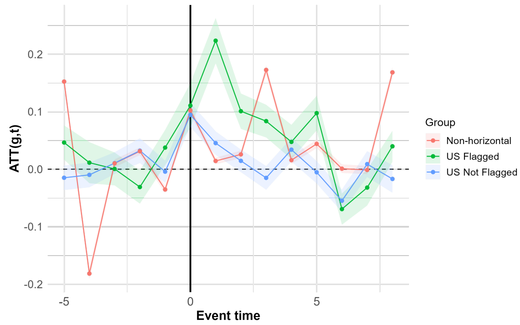

(b) App update

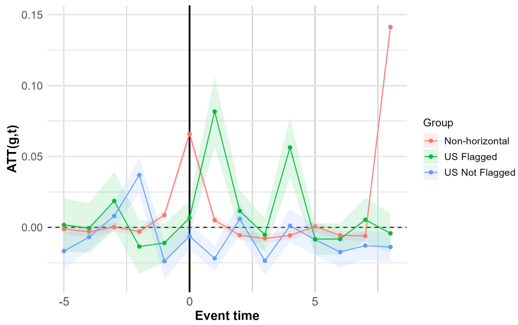

(c) Major update

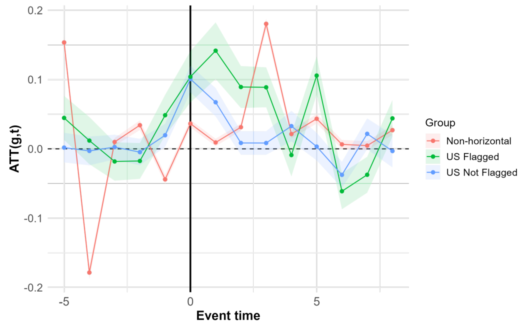

(d) Minor update

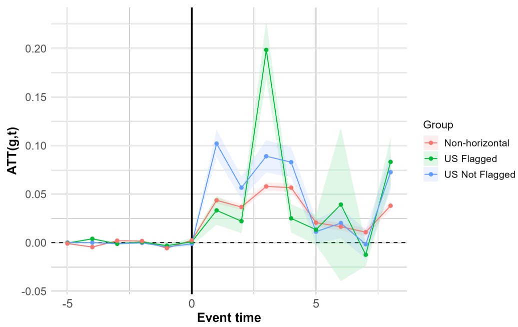

(e) Dismissed following period

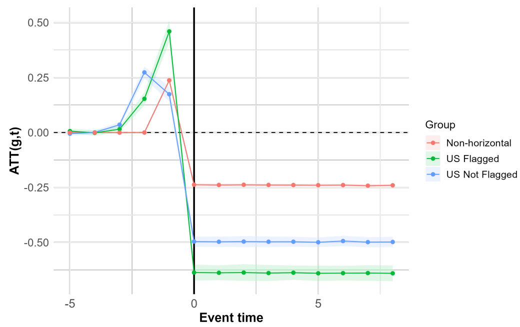

(f) Sold following period

Figure 4.4: Dynamic effects of M&A on app development by US policy screen

Notes: These figures show the event study results for mergers, capturing how various app-related metrics evolve before and after the merger event. The x-axis represents date relative to the merger, with the six week interval after the merger set at 0. Each sub-figure reflects a different app-related outcome, providing insights into the short- and long-term effects of mergers on app developer behaviour.

The event-study estimates in Figure 4.4 show clear heterogeneity by policy group. For US-flagged horizontal mergers, update rates pick up around the event, while app dismissals rise sharply and remain elevated-consistent with portfolio consolidation rather than expansion. Non-flagged horizontals display much smaller or short-lived movements across outcomes. By contrast, non-horizontal deals sustain higher overall and Major update activity, alongside modest price increases, and only a smaller increase in exits relative to flagged horizontals. Taken together, these dynamics match the ATT patterns in Table 4.7: concentration-screened mergers mainly reallocate effort within existing portfolios, whereas cross-market integrations tend to support incremental innovation and monetization.

4.6 Conclusion

This chapter develops a text-based framework for defining relevant markets and applies it to study how mergers and acquisitions affect non-price conduct in the app economy. Using text-defined markets constructed from app descriptions and update notes, the analysis identifies 1,914 horizontal overlaps, of which 375 are flagged by at least one policy screen and 84.5% by both the EU and US criteria. The vast majority of observed transactions are non-horizontal (18,727, about 91%), indicating that most acquisitions combine complementary rather than directly competing products.

Among horizontal mergers, those flagged by concentration screens exhibit clear signs of portfolio rationalization. Short-run app exits increase by roughly six percentage points, while update activity rises modestly, with no statistically or economically detectable change in listed prices. Horizontal mergers below policy thresholds display smaller pruning-around four percentage points-and little adjustment on update margins. By contrast, non-horizontal integrations are associated with higher update rates (about five percentage points), modest increases in listed prices (about \$0.12), and selective increases in app exits, consistent with active portfolio reallocation following scope expansion.

Event-study evidence supports this interpretation. In flagged horizontal cases, exit rates rise sharply at the time of the merger and remain elevated, while update activity increases briefly before stabilizing. These dynamics are consistent with post-merger consolidation and a reallocation of maintenance effort across surviving products rather than with price-based exploitation of increased market power. Horizontal mergers below the thresholds show weaker and less persistent responses. Non-horizontal mergers, in contrast, sustain higher update activity and modest price gains, reflecting integration across complementary product lines rather than rationalization within narrowly defined markets.

The observed price effects should be interpreted with care. Prices are identified only for the subset of paid apps, which constitute a small minority of the overall app population. Accordingly, the estimated price increases for non-horizontal mergers reflect changes among paid apps and do not imply broad-based price effects at the platform level. For the majority of apps-those offered at zero price-post-merger adjustment occurs primarily along non-price margins.

Robustness checks indicate that these patterns are not driven by specific clustering choices or market partitions. Overlap assignments and screening rates vary little across algorithms (HDBSCAN versus K-means) and levels of granularity (300–1,200 app caps), and cluster stability metrics (ARI and Jaccard similarity) confirm consistent results. The findings hold under both EU and US screening criteria and are robust to alternative constructions of concentration at the app and developer level.

The policy interpretation of these results is descriptive. Current merger screens identify roughly one-fifth of horizontal overlaps-the same subset that exhibits the strongest portfolio rationalization-while most below-threshold cases show more limited adjustment, and non-horizontal deals adjust primarily outside the screened markets. These findings suggest that combining text-based market definition with observable non-price indicators, such as exit and update activity, can help characterize merger effects in settings where price-based tests are less informative.

Overall, the chapter offers both a methodological and a substantive contribution. Methodologically, it provides a scalable approach to market delineation based on product text where price data are scarce. Substantively, it documents how consolidation in digital markets primarily reshapes portfolio composition and maintenance activity, while price effects-when present-are confined to paid apps and arise mainly in non-horizontal integrations.

Future work and limitations. Future extensions can move beyond outcomes for acquirers and targets to examine how rival apps within the same text-defined market adjust development effort, entry, and exit following a merger, allowing an assessment of whether consolidation dampens or stimulates market-level competitive activity. A related direction is to study changes in overall product variety and entry rates within clusters, linking observed post-merger rationalization to broader shifts in market dynamism. Further work could also trace the post-acquisition trajectory of acquired apps-whether they are actively developed, maintained at low intensity, or quietly discontinued-and assess whether some horizontal acquisitions function as *killer acquisitions* by suppressing future development of overlapping products.

An additional extension would be to adapt the text-based market framework to study complementarities and potential efficiencies in vertical or non-horizontal mergers. This could be achieved using nested or hierarchical clustering schemes, in which upstream and downstream product spaces are identified separately and linked through textual similarity or functional overlap. Such an approach could help distinguish efficiency-enhancing integration from consolidation-driven rationalization in non-horizontal transactions. Progress along these dimensions will require richer product-level data, such as feature metadata or usage measures, and improved within-developer control groups to capture heterogeneous responses.

4.A Appendix

4.A.1 Stability/overlap metrics (ARI, Jaccard)

To evaluate the robustness of our clustering choices, we report two widely used measures of overlap: the ARI and the Jaccard index as shown in Figure 4.5. Both compare how similar two alternative partitions of the same set of apps are, relative to the baseline specification. The ARI adjusts for agreement that would occur purely by chance, so values close to 1 indicate very high stability of cluster assignments, while values near 0 suggest similarity no better than random. The Jaccard index instead measures the proportion of app pairs that are grouped together in both specifications out of all pairs that are grouped together in at least one, again ranging from 0 (no overlap) to 1 (perfect overlap). In our context, higher ARI and Jaccard scores mean that 'text-defined markets' remain largely unchanged when we vary the clustering algorithm or cap, providing reassurance that the chosen baseline market definition is not an artifact of one particular specification.

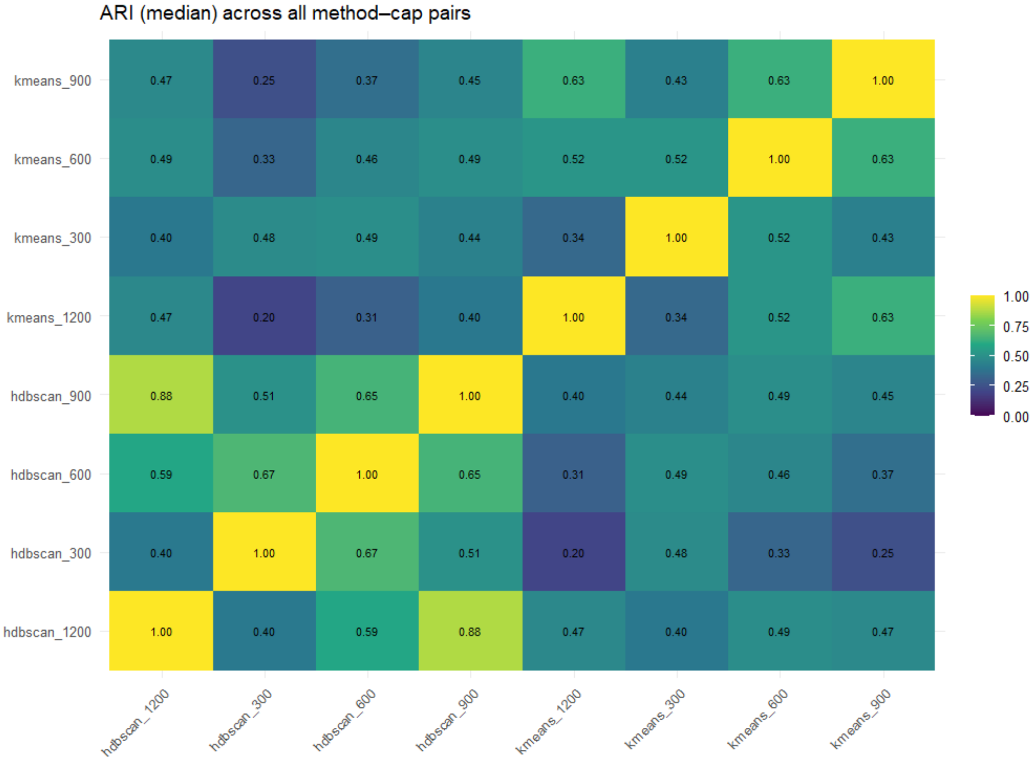

(a) ARI overlap across clustering specifications.

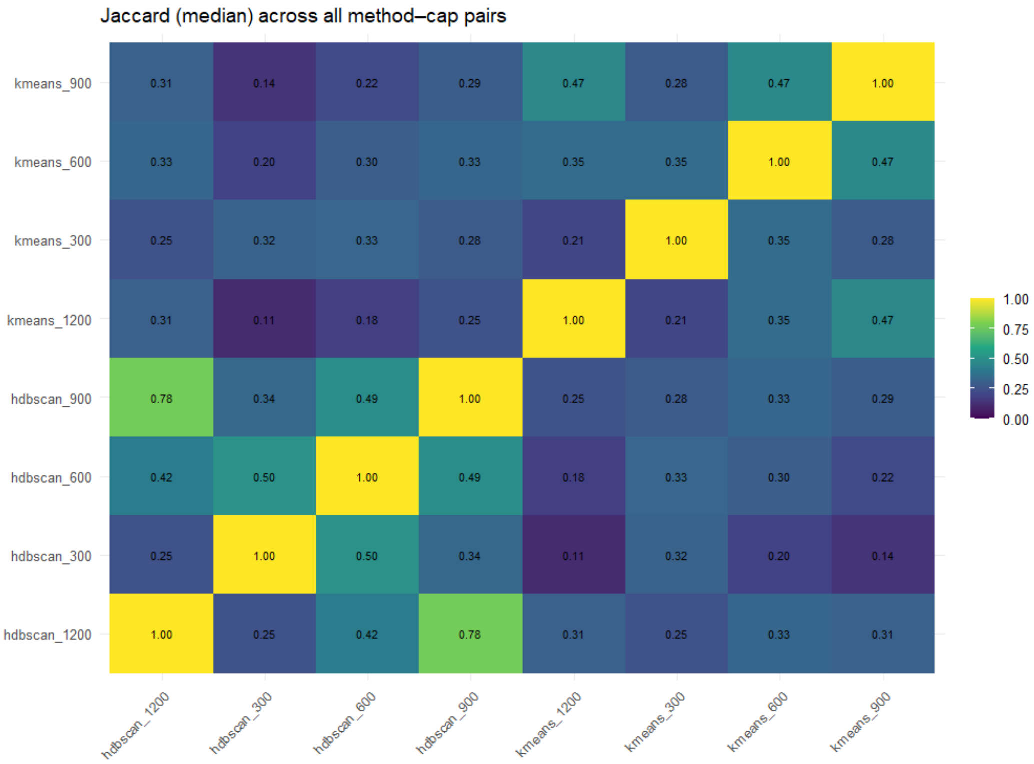

(b) Jaccard similarity overlap across clustering specifications.

Figure 4.5: Stability of cluster assignments across alternative caps and methods, measured by (a) ARI and (b) Jaccard similarity.

4.A.2 Extended market construction (text-defined markets)

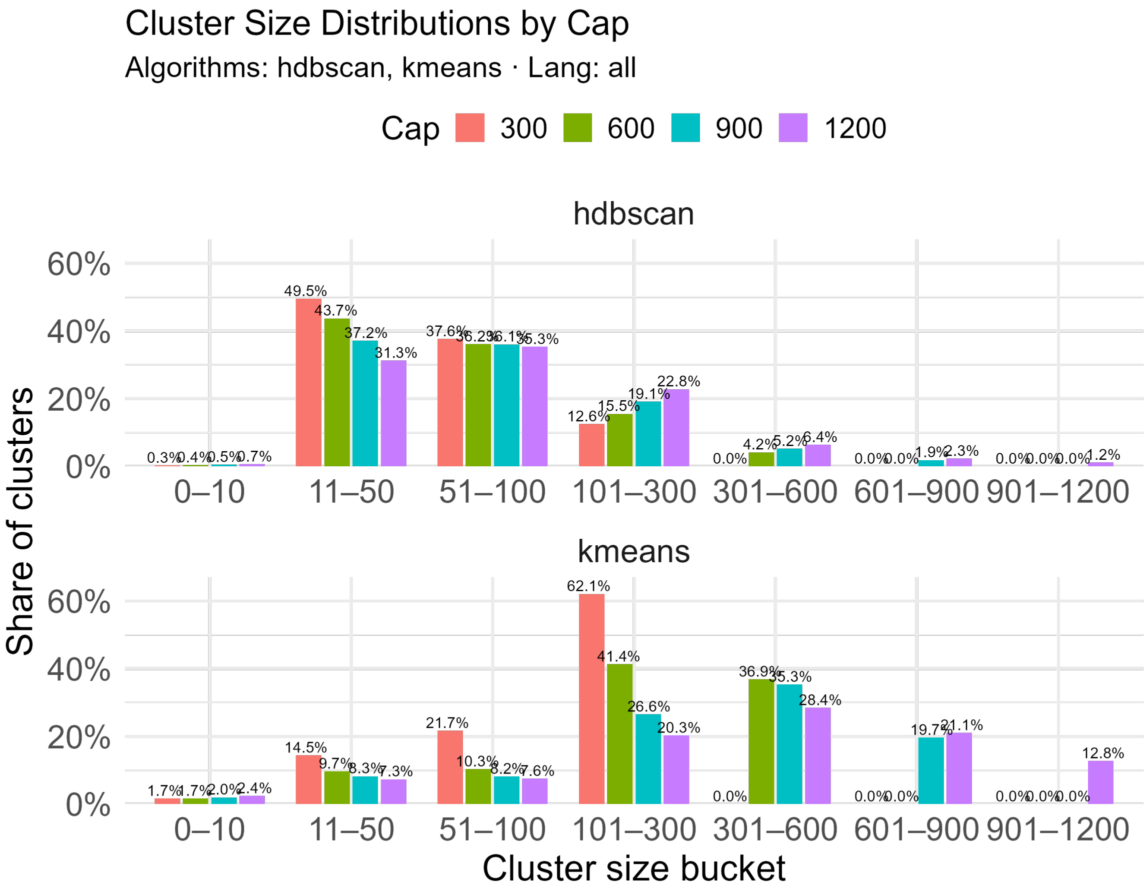

Figure 4.6: Cluster size distributions under alternative caps (300, 600, 900, 1200)

Table 4.8: Extended Baseline Market Properties by Language, Cap, and Method

	English (en)						Non-English (nonen)					
Method–Cap	Wgt. Mean Size	Unwgt. Mean Size	Min Size	Max Size	Small (1–10 apps)	HHI (Q1, Q3)	Wgt. Mean Size	Unwgt. Mean Size	Min Size	Max Size	Small (1–10 apps)	HHI (Q1, Q3)
HDBSCAN–300	64.38	64.30	20	300	0	(4.9, 25.8)	54.21	46.18	1	299	402 (7+4+131+260)	(60.3, 423.4)
HDBSCAN–600	84.60	89.10	20	600	0	(10.0, 42.8)	72.05	61.33	1	596	402	(122.9, 628.1)
HDBSCAN–900	107.27	122.87	20	900	0	(14.7, 99.9)	88.23	72.31	1	900	402	(261.7, 698.1)
HDBSCAN–1200	132.56	152.11	20	1199	0	(19.6, 150.8)	105.00	82.84	1	1179	402	(302.5, 718.2)
K-means–300	136.66	133.43	6	300	511 (0+0+0+511)	(8.7, 63.7)	101.47	75.22	1	300	487 (7+4+131+345)	(123.9, 556.6)
K-means–600	262.99	255.71	6	600	125 (0+0+0+125)	(17.6, 128.1)	160.53	112.40	1	700	407 (7+4+131+265)	(235.5, 695.1)
K-means–900	387.64	378.16	6	900	33	(26.1, 187.5)	193.07	134.83	1	900	404 (7+4+131+262)	(357.8, 699.2)
K-means–1200	509.34	505.83	8	1200	10	(34.2, 242.3)	215.73	153.76	1	1198	402 (7+4+131+260)	(385.2, 744.2)

Notes: The table reports extended diagnostics for English and non-English apps across clustering methods (HDBSCAN, K-means) and caps (300–1200). Weighted and unweighted mean size are average apps per cluster. Min and Max report extreme cluster sizes. 'Small (1–10 apps)' aggregates the number of very small markets, with parentheses showing the breakdown (singletons, two-app, three–four, five–ten apps). HHI (Q1, Q3) reports interquartile ranges of the HHI across Categories (0–10,000 scale). These statistics confirm predictable shifts: larger caps yield bigger and more concentrated clusters; non-English markets are systematically smaller and more prone to very small clusters, while also exhibiting higher concentration levels.

4.A.3 Defining markets: alternative clusterings, same conclusions

Table 4.9: Recovery of baseline policy flags Under alternative clusterings

		EU buckets		
Comparison group	Regime	Flagged	Unflagged	Undefined
Panel A. EU screens				
HDBSCAN vs K-means (same caps)	EU	0.35	0.68	0.86
Within HDBSCAN (other caps)	EU	0.61	0.63	0.74
Panel B. US screens				
HDBSCAN vs K-means (same caps)	US	0.36	0.68	0.86
Within HDBSCAN (other caps)	US	0.62	0.62	0.74

Notes: Baseline specification is HDBSCAN with *caps* = 900. Entries report the average *recall* of the baseline policy bucket when the same apps are classified with an alternative clustering. Recall is defined as the share of apps flagged (or unflagged, or undefined) under the baseline that receive the same bucket under the comparison specification. 'Within HDBSCAN' compares HDBSCAN(900) to HDBSCAN(300, 600, 1200); 'HDBSCAN vs K-means (same caps)' compares HDBSCAN(900) to K-means(900). High values suggest the policy screen is stable to the change in clustering. The 'Undefined' bucket naturally shows high recall because it pools apps without sufficient market information.

Table 4.10: HHI-based policy screens under K-means clustering

	300		600		900		1200	
Panel A. Non-horizontal M&A (no overlapping market)	18,651	90.7	17,508	90.1	16,823	89.6	16,343	89.4
Panel B. Horizontal M&A - policy screening within overlapping markets								
EU competition policy								
Flagged	329	17.2	274	14.2	251	12.9	237	12.2
Not flagged	1,589	82.8	1,655	85.8	1,701	87.1	1,702	87.8
Total (EU horizontal)	1,918	100.0	1,929	100.0	1,952	100.0	1,939	100.0
US (DOJ/FTC structural-presumption test)								
Flagged	367	19.9	306	16.8	261	16	251	12.9
Not flagged	1,551	80.1	1,623	84.0	1,691	86.6	1,688	87.1
Total (US horizontal)	1,918	100.0	1,929	100.0	1,952	100.0	1,939	100.0
Panel C. Overlap composition among screened-in								
EU flagged (only)	16	4.2	19	5.8	34	11.5	29	10.4
US flagged (only)	54	14.1	51	15.7	44	14.9	43	15.4
Both (EU ∩ US)	313	81.7	255	78.5	217	73.6	208	74.2
Total flagged (EU ∪ US)	383	100.0	325	100.0	295	100.0	280	100.0

Notes: This table reports developer-level merger classifications under alternative K-means market granularities (*caps* $= 300, 600, 900, 1200$). Panel A lists **non-horizontal** M&A, defined as transactions between parties without overlapping market presence. Panel B focuses on **horizontal** cases-where merging developers share the same relevant market-and applies two regulatory screens. EU 'Flagged' corresponds to deals *outside the EU safe harbours*: (i) $1000 \leq HHI_{\text{post}} < 2000$ and $\Delta HHI \geq 250$, or (ii) $HHI_{\text{post}} \geq 2000$ and $\Delta HHI \geq 150$. US 'Flagged' corresponds to the structural-presumption test, $HHI_{\text{post}} \geq 1800$ and $\Delta HHI \geq 100$. Percentages in Panel B are computed **within the horizontal subgroup** of each jurisdiction. Panel C shows the overlap composition among flagged events, distinguishing EU-only, US-only, and jointly flagged clusters. All shares refer to cluster-level counts..

Table 4.11: App composition under K-means clustering (developer-level markets)

	300			600			900			1200		
	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total
Panel A. Non-horizontal M&A (no overlapping market)	23,612	1,102,214	1,125,826	23,142	2,047,938	2,071,080	22,769	2,922,445	2,945,214	22,618	3,738,680	3,761,298
Panel B. Horizontal M&A - policy screening within overlapping markets												
EU competition policy												
Flagged (screened-in)	1,026	17,991	19,017	1,205	25,460	26,665	1,005	38,338	39,343	758	45,277	46,035
Not flagged (safe harbours)	2,942	95,402	98,344	3,228	200,204	203,432	3,801	305,453	309,254	4,199	397,762	401,961
Total (EU horizontal)	3,968	113,393	117,361	4,433	225,664	230,097	4,806	343,791	348,597	4,957	443,039	447,996
US (DOJ/FTC structural-presumption test)												
Flagged	1,083	20,079	21,162	1,388	29,554	30,942	1,099	39,651	40,750	776	48,918	49,694
Not flagged (below thresholds)	2,885	93,314	96,199	3,045	196,110	199,155	3,707	304,140	307,847	4,181	394,121	398,302
Total (US horizontal)	3,968	113,393	117,361	4,433	225,664	230,097	4,806	343,791	348,597	4,957	443,039	447,996

Notes: 'Parties' counts distinct apps owned by the merging developers at the strict last pre-t_0 snapshot; 'Cluster total' counts all apps in the corresponding market; 'Rivals' equals Cluster total − Parties. Panel A uses the no-overlap group (non-horizontal). EU 'Flagged' corresponds to deals *outside the EU safe harbours*: (i) $1000 \leq HHI_{\text{post}} < 2000$ and $\Delta HHI \geq 250$, or (ii) $HHI_{\text{post}} \geq 2000$ and $\Delta HHI \geq 150$. US 'Flagged' corresponds to the structural-presumption test, $HHI_{\text{post}} \geq 1800$ and $\Delta HHI \geq 100$. All figures are based on cluster-level counts of distinct apps.

Table 4.12: HHI-based policy screens under HDBSCAN clustering

	300		600		900		1200	
Panel A. Non-horizontal M&A (no overlapping market)	19,985	91.4	19,237	91.0	18,727	90.7	18,278	90.6
Panel B. Horizontal M&A - policy screening within overlapping markets								
EU competition policy								
Flagged	384	20.3	357	18.8	333	17.4	317	16.6
Not flagged	1,507	79.7	1,545	81.2	1,581	82.6	1,590	83.4
Total (EU horizontal)	1,891	100.0	1,902	100.0	1,914	100.0	1,907	100.0
US (DOJ/FTC structural-presumption test)								
Flagged	427	22.6	391	20.6	359	18.8	344	19.0
Not flagged	1,464	77.4	1,511	79.4	1,555	81.2	1,466	81.0
Total (US horizontal)	1,891	100.0	1,902	100.0	1,914	100.0	1,810	100.0
Panel C. Overlap composition among screened-in								
EU flagged (only)	6	1.4	9	2.2	16	4.3	17	4.7
US flagged (only)	49	11.3	43	10.8	42	11.2	44	12.2
Both (EU $\cap$ US)	378	87.3	348	87.0	317	84.5	300	83.1
Total flagged (EU $\cup$ US)	433	100.0	400	100.0	375	100.0	361	100.0

Notes: This table reports developer-level merger classifications under alternative **HDBSCAN** market granularities (*caps* $= 300, 600, 900, 1200$). Panel A lists **non-horizontal** M&A, defined as transactions between developers without overlapping market presence. Panel B restricts to **horizontal** cases-where overlap exists-and applies policy thresholds used by the European Commission and the US Department of Justice/Federal Trade Commission. For the EU, a merger is *Flagged* if it lies *outside the safe harbours*: (i) $1000 \leq HHI_{\text{post}} < 2000$ and $\Delta HHI \geq 250$, or (ii) $HHI_{\text{post}} \geq 2000$ and $\Delta HHI \geq 150$. For the US, a merger is *Flagged* if it meets the structural-presumption thresholds, $HHI_{\text{post}} \geq 1800$ and $\Delta HHI \geq 100$. Percentages in Panel B are computed **within the horizontal subgroup** of each jurisdiction. Panel C decomposes the overlap among flagged clusters, showing EU-only, US-only, and jointly flagged cases. All shares refer to cluster-level counts. Presentation and terminology follow standard competition-policy conventions (OECD, EC Merger Guidelines, and DOJ/FTC Horizontal Merger Guidelines).

Table 4.13: App composition under HDBSCAN clustering (developer-level markets)

	300			600			900			1200		
	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total	Parties	Rivals	Cluster total
Panel A. Non-horizontal M&A (no overlapping market)	24,171	619,271	643,442	23,767	1,076,114	1,099,881	23,646	1,538,573	1,562,219	23,356	1,977,297	2,000,653
Panel B. Horizontal M&A - policy screening within overlapping markets												
EU competition policy												
Flagged (screened-in)	1,108	10,127	11,235	1,235	16,493	17,728	1,057	23,776	24,833	1,041	28,958	29,999
Not flagged (safe harbours)	2,350	50,792	53,079	2,575	92,205	94,780	2,926	147,748	150,674	3,183	193,459	196,642
Total (EU horizontal)	3,458	60,856	64,314	3,810	108,698	112,508	3,983	171,524	175,507	4,224	222,417	226,641
US (DOJ/FTC structural-presumption test)												
Flagged	1,176	11,587	12,763	1,298	17,610	18,908	1,103	24,255	25,258	1,145	30,681	31,826
Not flagged (below thresholds)	2,282	51,551	49,269	2,512	91,088	93,600	2,880	147,369	150,249	3,079	191,736	194,815
Total (US horizontal)	3,458	60,856	64,314	3,810	108,698	112,508	3,983	171,524	175,507	4,224	222,417	226,641

Notes: 'Parties' counts distinct apps owned by the merging developers at the strict last pre-t_0 snapshot; 'Cluster total' counts all apps in the corresponding market; 'Rivals' equals Cluster total − Parties. Panel A uses the no-overlap group (non-horizontal). EU 'Flagged' corresponds to deals outside the EU safe harbours; US 'Flagged' corresponds to the structural-presumption test ($HHI \geq 1800$ and $\Delta HHI \geq 100$). All figures are based on cluster-level counts of distinct apps. Minor discrepancies in totals across caps reflect filtering of multi-deal developers and rounding of percentages.

Across market-definition choices, the vast majority of *event–cluster classifications* are non-horizontal: approximately 90–91% of observations fall into the 'no overlapping market' category under both K-means and HDBSCAN, corresponding to about 16.3k–20.0k cluster-level observations depending on granularity (see Table 4.10 and Table 4.12). Within the horizontal subset, screening incidence declines as markets are partitioned more finely. Under K-means, the EU 'screened-in' rate decreases from 17.2% at *cap*=300 to 12.2% at *cap*=1200, while the US structural-presumption rate falls from 20.0% to 12.9%. HDBSCAN yields uniformly higher screening rates at comparable granularities-EU: 20.3% to 16.6%; US: 22.6% to 19.0%-suggesting denser, demand-consistent clusters expose more overlap than equal-sized K-means partitions. Among flagged cases, joint EU∩US triggers dominate: under K-means they account for 81.7% at *cap*=300 and 74.2% at *cap*=1200; under HDBSCAN the corresponding shares are 87.3% and 83.1%. EU-only flags, while small, rise with granularity (K-means: 4.2% → 10.4%; HDBSCAN: 1.4% → 4.7%), and US-only remains in a narrow 11–15% band across definitions.

These patterns are consistent with standard concentration mechanics and the construction of text-defined markets. Finer partitions mechanically attenuate HHI and ΔHHI by allocating products to narrower niches, reducing the probability that merging parties share a market and, conditional on overlap, lowering combined shares. Conversely, the density-based HDBSCAN definition tracks 'functional' product space where user demand bunches, yielding higher overlap and more screens than equal-size K-means. The predominance of joint EU∩US flags indicates that most at-risk deals occur where both overall concentration and incremental concentration are sufficiently large to surpass both jurisdictions' thresholds, while the rise in EU-only flags at finer granularities reflects the EU's outside–safe-harbours logic that can be triggered when either HHI or ΔHHI is large in tightly defined niches, even if the US pair of thresholds is not jointly met (see Table 4.10, Table 4.12).

4.A.4 Extended empirical result

Table 4.14: ATT by EU policy screen

	EU Flagged	EU Not Flagged	Non-horizontal
App outcomes			
(1)Price ($)	0.014 (0.028)	0.063*** (0.022)	0.117*** (0.017)
(2)App update	0.013 (0.011)	-0.002 (0.009)	0.047*** (0.003)
(2a)Major update	0.009* (0.006)	-0.012*** (0.004)	0.014*** (0.001)
(2b)Minor update	0.004 (0.011)	0.011 (0.008)	0.032*** (0.003)
(3)Dismissed following period	0.055*** (0.009)	0.040*** (0.002)	0.035*** (0.001)
(4)Sold following period	-0.648*** (0.000)	-0.488*** (0.000)	-0.222*** (0.000)
Time FE	Yes	Yes	Yes
App FE	Yes	Yes	Yes
Apps (treated), N_{tr}	700	1,894	16,544
Apps (controls), N_{ctrl}	19,564	19,564	19,564
Observations (N_{obs})	352,462	373,147	676,667

Notes: Entries report overall ATT from the staggered DiD estimator of Callaway and Sant'Anna (2021), computed separately for three EU policy groups defined at the event time: *EU Flagged*, *EU Not flagged*, and *Non-horizontal*. Each cell shows the point estimate with clustered standard errors in parentheses; standard errors are clustered at the app level. All specifications include time and app fixed effects. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. N_{obs} reports the modal (majority) number of app–time observations across outcomes in each policy group. For outcomes with minor additional missingness, N_{obs} is slightly lower. For Non-horizontal specifically: 'Dismissed' N_{obs} = 664,944 and 'Sold' N_{obs} = 642,222 (vs. 676,667 for price/update/key/small). Analogous small reductions occur, if any, in other groups.

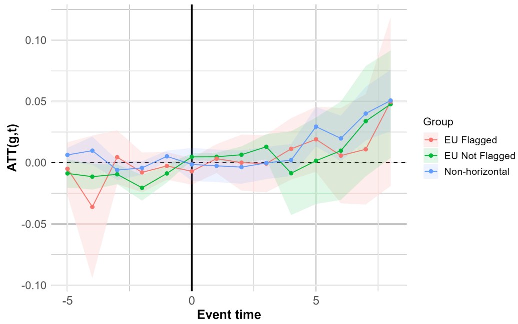

(a) Price ($)

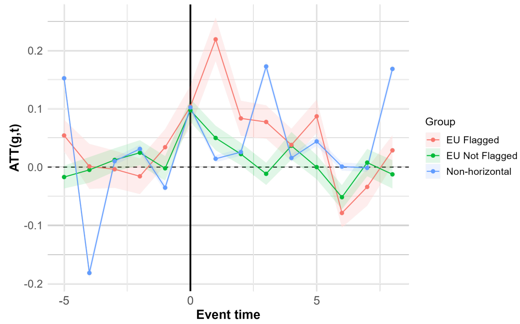

(b) App update

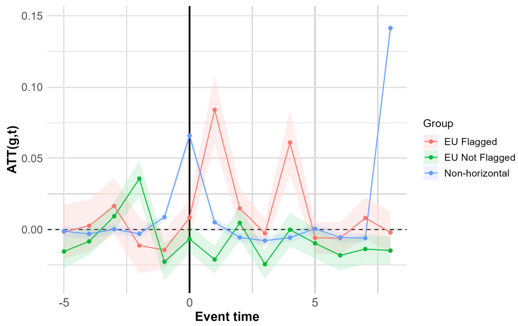

(c) Major update

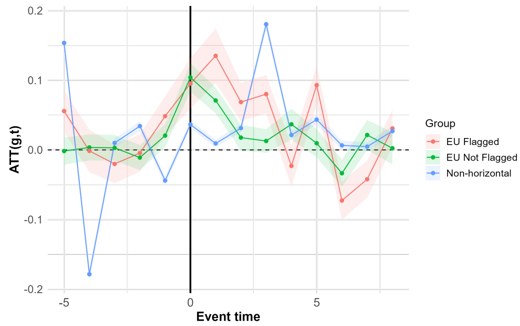

(d) Minor update

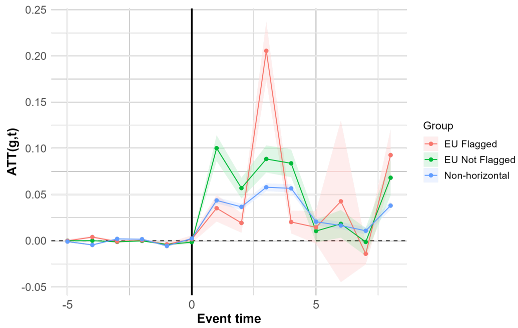

(e) Dismissed following period

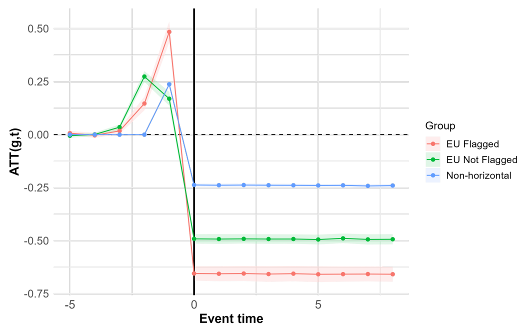

(f) Sold following period

Figure 4.7: Comparative event study: impact of M&A on app development

Notes: These figures show the event study results for mergers, capturing how various app-related metrics evolve before and after the merger event. The x-axis represents date relative to the merger, with the six week interval after the merger set at 0. Each sub-figure reflects a different app-related outcome, providing insights into the short- and long-term effects of mergers on app developer behaviour.

Chapter 5

Conclusion

This thesis has analysed how mergers and acquisitions reshape competition in digital markets where platforms intermediate product discovery and quality signals. Across three chapters, it has focused on environments in which conventional price-based merger tools provide an incomplete account of competitive effects, either because prices are only one of several strategic variables or because user-side prices are zero. Taken together, the essays show that digital consolidation often operates through adjustments in perceived quality, innovation, and portfolio composition rather than through large, observable price movements, and they develop tools that render these channels empirically tractable for merger assessment.

Chapter 2 provides a structural perspective on platform-driven quality contributions in a non-horizontal merger. It develops a duopoly framework in which a dominant platform allocates 'quality contributions'—for example via ratings or visibility—across competing sellers and embeds this mechanism in a RCL model of the US robotic vacuum market. Using the Amazon–iRobot case, the chapter shows that modest, algorithmically driven boosts to the acquired brand's ratings can materially reallocate market shares, expand the market, raise platform and firm profits, and shift incentives to invest in product quality. This highlights a non-price channel through which vertical integration in digital markets affects competition via control of reputational and informational infrastructure rather than direct price changes.

Chapter 3 turns to a broad cross-section of M&A among app developers on Google Play and provides a reduced-form analysis of non-price firm behaviour within a single digital ecosystem. Using panel data on roughly 8 million apps, 1.8 million developers, and 1,387 US M&A, the chapter traces how portfolios and development strategies evolve around merger events, distinguishing between mergers and acquisitions and between horizontal and non-horizontal deals. Consolidation reduces new app launches and external purchases while limiting removals and sales, leading to a net contraction in product variety but increased maintenance and quality-enhancing updates to existing apps. Acquisitions feature a more pronounced contraction in new launches than mergers, and horizontal and non-horizontal events display distinct patterns of portfolio reallocation. These patterns underscore that entry, variety, and innovation respond meaningfully to digital mergers even when posted prices remain relatively stable.

Chapter 4 tackles the complementary problem of market definition and screening in settings where user-side prices are frequently zero. It develops a framework for mapping products into 'text-defined markets' using app descriptions and update notes, combining sentence embeddings and clustering with targeted validation. These text-defined markets support standard concentration measures and can be linked to staggered DiD estimators. Applied to Google Play (2015–2019), the framework maps 1,914 horizontal M&A events into markets and screens them using EU- and US-style concentration thresholds. Flagged deals exhibit sharper portfolio rationalisation —higher exit and modestly increased update activity—but no clear price effects, while below-threshold horizontal and non-horizontal deals primarily reallocate update effort and generate only small price changes. The chapter shows that text-based approaches can deliver operational market definitions and policy-relevant concentration screens in data-rich but price-poor environments, providing an applied template for merger analysis when SSNIP-style tools are ill-posed.

Viewed jointly, the three chapters provide a coherent account of non-price merger effects in digital markets. Chapter 2 shows how platform control over quality signals can amplify the competitive impact of vertical integration. Chapter 3 documents, at scale, how consolidation among independent developers shifts effort from extensive-margin expansion towards intensive-margin maintenance and upgrading. Chapter 4 demonstrates that market definitions recovered from product text can underpin concentration-based screens and event-style evaluations of non-price outcomes. Together, these contributions move beyond a purely price-centric view of merger control, emphasising the role of algorithmically mediated quality signals, portfolio decisions, and text-based market boundaries in shaping competition on digital platforms.

The analysis in this thesis has limitations that delineate avenues for further research. The structural model in Chapter 2 abstracts from richer dynamics such as reputation accumulation, learning, and multi-homing across platforms. Extending the framework to incorporate dynamic rating processes or multi-platform interactions would yield a more complete picture of how quality contributions evolve over time and across ecosystems. The reduced-form designs in Chapters 3 and 4, while based on modern event-study and staggered DiD methods, inevitably rely on assumptions about parallel trends and the absence of unobserved shocks correlated with treatment. Future work could combine structural and design-based approaches, for example by using text-defined markets as inputs into structural demand-supply models or by embedding richer dynamic models of entry, exit, and innovation into the app-economy setting.

The empirical focus is also deliberately narrow in sectoral terms, concentrating on robotic vacuum cleaners and mobile applications. While these are salient and policy-relevant digital markets, external validity beyond these settings is not automatic. Applying the quality-contribution framework and text-defined market construction to other environments—such as video streaming, search, cloud services, or digital advertising—would help assess the extent to which the mechanisms documented here generalise across different forms of platform intermediation and business models. Relatedly, the non-price outcomes studied in this thesis centre on innovation, variety, and prices; complementary work could examine privacy, data practices, labour outcomes, and content diversity as additional margins along which digital mergers may have substantial welfare effects.

Finally, the results have implications for how competition authorities might prioritise and design digital merger review. The analysis of platform-driven quality contributions suggests that agencies should pay closer attention to how vertical integration interacts with control over ranking, recommendations, and ratings, and to whether commitments or remedies targeting algorithmic governance are warranted in specific cases. The evidence on portfolio adjustment and innovation highlights the need to look beyond static price effects when assessing welfare in markets where product variety and incremental updates are central to consumer benefits. The text-defined market framework shows that NLP tools can be used to construct economically meaningful markets and to generate concentration-based screens even when prices are zero, helping to focus enforcement on cases with the greatest potential for harmful non-price effects.

In sum, this thesis contributes to an emerging research programme that views digital mergers as operating through information, quality, and innovation as much as through prices. By combining structural modelling, large-scale empirical analysis, and NLP-based market construction, the three chapters provide building blocks for understanding and regulating consolidation in algorithmically mediated markets. Further work that broadens sectoral coverage, refines models of platform behaviour, and extends the range of non-price outcomes will be important for keeping merger policy effective in digital economies where data, visibility, and quality signals are central to competition.

Bibliography

Affeldt, Pauline and Reinhold Kesler, "Competitors' Reactions to Big Tech Acquisitions: Evidence from Mobile Apps," Technical Report 2021.

Aghion, Philippe, Roland Bénabou, Ralf Martin, and Alexandra Roulet, "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?," *American Economic Review: Insights*, 2023, *5* (1), 1–20.

Amit Gandhi and Jean-Francois Houde, "Measuring Substitution Patterns in Differentiated-Products Industries," *NBER woking paper*, 2019.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira, "Patents to products: Product innovation and firm dynamics," 2020.

Argentesi, Elena, Paolo Buccirossi, Emilio Calvano, Tomaso Duso, Alessia Marrazzo, and Salvatore Nava, "Merger Policy in Digital Markets: An Ex Post Assessment," *Journal of Competition Law and Economics*, 3 2021, *17* (1), 95–140.

Armstrong, Mark, "Competition in two-sided markets," *The RAND Journal of Economics*, 9 2006, *37* (3), 668–691.

_ **and Jidong Zhou**, "Paying for Prominence," *The Economic Journal*, 11 2011, *121* (556), F368–F395.

Arrow, K J, "Economic Welfare and the Allocation of Resources for Invention," in Charles K Rowley, ed., *Readings in Industrial Economics: Volume Two: Private Enterprise and State Intervention*, London: Macmillan Education UK, 1972, pp. 219–236.

Ashenfelter, Orley and Daniel Hosken, "The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin," *https://doi.org/10.1086/605092*, 8 2010, *53* (3), 417–466.

Ashenfelter, Orley C, Daniel S Hosken, and Matthew C Weinberg, "The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool," *American Economic Journal: Economic Policy*, 2013, *5* (1), 239–261.

Azar, José, Martin C. Schmalz, and Isabel Tecu, "Anticompetitive Effects of Common Ownership," *Journal of Finance*, 8 2018, *73* (4), 1513–1565.

Bennato, Anna Rita, Stephen Davies, Franco Mariuzzo, and Peter Ormosi, "Mergers and innovation: Evidence from the hard disk drive market," *International Journal of Industrial Organization*, 6 2021, *77*.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," Technical Report 4 1995.

Berry, Steven T and Joel Waldfogel, "DO MERGERS INCREASE PRODUCT VARIETY? EVIDENCE FROM RADIO BROADCASTING*," Technical Report 2001.

Blonigen, Bruce and Justin Pierce, "Evidence for the Effects of Mergers on Market Power and Efficiency," Technical Report, Cambridge, MA 2016.

Bourreau, Marc, Bruno Jullien, and Yassine Lefouili, "Horizontal Mergers and Incremental Innovation," 2021.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "1. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *An Agenda*, University of Chicago Press, 2019, pp. 23–60.

Cabral, Luís, "Big Tech acquisitions," *International Journal of Industrial Organization*, December 2025, *103*, 103127.

Caillaud, Bernard and Bruno Jullien, "Chicken & egg: Competition among intermediation service providers," *RAND journal of Economics*, 2003, pp. 309–328.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 12 2021, *225* (2), 200–230.

Calvano, Emilio and Michele Polo, "Market power, competition and innovation in digital markets: A survey," *Information Economics and Policy*, 3 2021, *54*, 100853.

Cassiman, Bruno, Massimo G Colombo, Paola Garrone, and Reinhilde Veugelers, "The impact of M&A on the R&D process: An empirical analysis of the role of technological- and market-relatedness," *Research Policy*, 2005, *34* (2), 195–220.

Chandra, Ambarish and Allan Collard-Wexler, "Mergers in Two-Sided Markets: An Application to the Canadian Newspaper Industry," *Journal of Economics & Management Strategy*, 12 2009, *18* (4), 1045–1070.

Chevalier, Judith A and Dina Mayzlin, "The Effect of Word of Mouth on Sales: Online Book Reviews," *Journal of Marketing Research*, 2006, *43* (3), 345–354.

Chipty, Tasneem, "Vertical Integration, Market Foreclosure, and Consumer Welfare in the Cable Television Industry," *American Economic Review*, 2001, *91* (3), 428–453.

Choi, Hana and Carl F. Mela, "Monetizing Online Marketplaces," *Marketing Science*, 11 2019, *38* (6), 948–972.

Commission, European, "Market Definition Notice (Revised)," Technical Report 2024.

Competition and Markets Authority, "Merger Assessment Guidelines," Technical Report 2021.

Condorelli, Daniele and Jorge Padilla, "Harnessing Platform Envelopment in the Digital World," *Journal of Competition Law & Economics*, 6 2020, *16* (2), 143–187.

Cowan, Simon, Peter Ormosi, and Joe Perkins, "New directions in competition policy: an overview," *Oxford Review of Economic Policy*, 01 2025, *40* (4), 687–695.

Crawford, Gregory S and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 2012, *102* (2), 643–685.

Crawford, Gregory S., Oleksandr Shcherbakov, and Matthew Shum, "Quality overprovision in cable television markets," 3 2019.

Cunningham, Colleen, Florian Ederer, and Song Ma, "Killer Acquisitions," *Journal of Political Economy*, 11 2020, *129* (3), 649–702.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 9 2020, *110* (9), 2964–2996.

de Cornière, Alexandre and Greg Taylor, "Integration and search engine bias," *The RAND Journal of Economics*, 9 2014, *45* (3), 576–597.

Deneckere, Raymond and Carl Davidson, "Incentives to Form Coalitions with Bertrand Competition," *The RAND Journal of Economics*, 1985, *16* (4), 473–486.

Etro, Federico, "Mergers of complements and entry in innovative industries," *International Journal of Industrial Organization*, 7 2019, *65*, 302–326.

European Commission, "Guidelines on the Assessment of Horizontal Mergers under the Council Regulation on the Control of Concentrations between Undertakings," February 2004. Official Journal C 31, 05/02/2004, pp. 5–18.

_ , "Competition Policy for the Digital Era," Report prepared by Special Advisers J. Crémer, Y.-A. de Montjoye, and H. Schweitzer 2019.

Evans, David S., "Antitrust Economics of Free," *SSRN Electronic Journal*, 4 2011.

Fan, Ying and Chenyu Yang, "American Economic Association Competition, Product Proliferation, and Welfare," *American Economic Journal: Microeconomics*, 2020, *12* (2), 99–134.

Farrell, Joseph and Carl Shapiro, "Horizontal Mergers: An Equilibrium Analysis," Technical Report 1 1990.

_ **and** _ , “Antitrust evaluation of horizontal mergers: An economic alternative to market definition,” *The B.E. Journal of Theoretical Economics*, March 2010, *10* (1).

Federal Trade Commission, “Complaint for Injunctive and Other Equitable Relief: FTC v. Amazon.com, Inc.,” Technical Report 2023.

_ , “FTC Sues Amazon for Illegally Maintaining Monopoly Power,” 2023.

Federico, Giulio, Gregor Langus, and Tommaso Valletti, “A simple model of mergers and innovation,” *Economics Letters*, 8 2017, *157*, 136–140.

_ **,** _ **, and** _ , “Horizontal mergers and product innovation,” *International Journal of Industrial Organization*, 7 2018, *59*, 1–23.

Fikru, Mahelet G and Luis Gautier, “Are big mergers welfare enhancing when there is environmental externality?,” *Energy Economics*, 2020, *87*, 104718.

Filistrucchi, Lapo, Damien Geradin, Eric Van Damme, and Pauline Affeldt, “MARKET DEFINITION IN TWO-SIDED MARKETS: THEORY AND PRACTICE,” *Journal of Competition Law & Economics*, 6 2014, *10* (2), 293–339.

Focarelli, Dario and Fabio Panetta, “Are Mergers Beneficial to Consumers? Evidence from the Market for Bank Deposits,” *American Economic Review*, 2003, *93* (4), 1152–1172.

for Competition European Commission, Directorate-General, “Antitrust: Commission sends Statement of Objections to Amazon for the use of non-public independent seller data and opens second investigation into its e-commerce business practices,” 2020.

Froeb, Luke M and Gregory J. Werden, “The Reverse Cellophane Fallacy in Market Delineation,” *Review of Industrial Organization*, 1992, *7* (2), 241–247.

Gal, Michal S and Daniel L Rubinfeld, “The hidden costs of free goods: Implications for antitrust enforcement,” *Antitrust Law Journal*, 2016, *80* (3), 521–562.

Gandhi, Ashvin, Brett Hollenbeck, and Zhijian Li, "Misinformation and Mistrust: The Equilibrium Effects of Fake Reviews on Amazon.com," Technical Report 2025.

Gautier, Axel and Joe Lamesch, "Mergers in the digital economy," *Information Economics and Policy*, 3 2021, *54*, 100890.

Gentzkow, Matthew and Jesse M Shapiro, "MEASURING GROUP DIFFERENCES IN HIGH-DIMENSIONAL CHOICES: METHOD AND APPLICATION TO CONGRESSIONAL SPEECH," *Econometrica*, 2019, *87* (4), 1307–1340.

Ghose, Anindya and Sang Pil Han, "Estimating demand for mobile applications in the new economy," *Management Science*, 2014, *60* (6), 1470–1488.

Gilbert, Richard J, "Competition, Mergers, and R&D Diversity," *Review of Industrial Organization*, 2019, *54* (3), 465–484.

_ **and Hillary Greene**, "Merging Innovation into Antitrust Agency Enforcement of the Clayton Act The FTC at 100: Centennial Commemorations and Proposals for Progress: Essays," *George Washington Law Review*, 2014, *83* (6), 1919–1947.

Gugler, Klaus, Florian Szucs, and Ulrich Wohak, "Using Natural Language Processing to Delineate Digital Markets," *Stan. Computational Antitrust*, 2024, *4*, 33.

Hagiu, Andrei and Bruno Jullien, "Why do intermediaries divert search?," *The RAND Journal of Economics*, 6 2011, *42* (2), 337–362.

Hall, Marshall and Nicolaus Tideman, "Measures of Concentration," *Journal of the American Statistical Association*, 1967, *62* (317), 162–168.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout, "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters," *Journal of Political Economy*, 3 2020, *128* (9), 3570–3625.

Haucap, Justus, Alexander Rasch, and Joel Stiebale, "How mergers affect innovation: Theory and evidence," *International Journal of Industrial Organization*, 3 2019, *63*, 283–325.

— **and Joel Stiebale**, "Non-price effects of mergers and acquisitions," Technical Report 402, Düsseldorf 2023.

Hoberg, Gerard and Gordon Phillips, "Text-Based Network Industries and Endogenous Product Differentiation," *Journal of Political Economy*, 8 2016, *124* (5), 1423–1465.

Hollenbeck, Brett, "Online reputation mechanisms and the decreasing value of chain affiliation," *Journal of Marketing Research*, 10 2018, *55* (5), 636–654.

Janßen, Rebecca, Reinhold Kesler, Michael E Kummer, Joel Waldfogel, Michael E Kummer EFRY, and Joel R Waldfogel Frederick Kappel Chair, "GDPR and the Lost Generation of Innovative Apps," Technical Report 2022.

Johnson, Justin P and Andrew Rhodes, "Multiproduct mergers and quality competition," *The RAND Journal of Economics*, 9 2021, *52* (3), 633–661.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun, "The Diffusion of New Technologies," *The Quarterly Journal of Economics*, 4 2025, *140* (2), 1299–1365.

Kaplow, Louis, "Why (Ever) Define Markets' (2010) 124(2) Harvard Law Review 437 MLA 9th ed. Kaplow, Louis," Technical Report 2 2010.

Katz, Michael L and Carl Shapiro, "Critical loss: Let's tell the whole story," *Antitrust*, 2002, *17*, 49.

Kenneth G. Elzinga and Thomas F. Hogarty, "The Problem of Geographic Market Delineation in Antimerger Suits," Technical Report 1973.

Kesler, Reinhold, Michael Kummer, and Patrick Schulte, "Competition and Privacy in Online Markets: Evidence from the Mobile App Industry," *SSRN Electronic Journal*, 2019.

Landes, William M and Richard A Posner, "Market power in antitrust cases," *J. Reprints Antitrust L. & Econ.*, 1997, *27*, 493.

Letina, Igor, Armin Schmutzler, and Regina Seibel, "KILLER ACQUISITIONS AND BEYOND: POLICY EFFECTS ON INNOVATION STRATEGIES," *International Economic Review*, 5 2024, *65* (2), 591–622.

Levin, Dan, "Horizontal Mergers: The 50-Percent Benchmark," *The American Economic Review*, 1990, *80* (5), 1238–1245.

Li, Sophia, Joe Mazur, Yongjoon Park, James Roberts, Andrew Sweeting, and Jun Zhang, "Repositioning and market power after airline mergers," *RAND Journal of Economics*, 3 2022, *53* (1), 166–199.

Liang, X, S Li, P Luo, and et al., "Green mergers and acquisitions and green innovation: an empirical study on heavily polluting enterprises," *Environmental Science and Pollution Research*, 2022, *29*, 48937–48952.

Long, Fei and Yunchuan Liu, "Platform Manipulation in Online Retail Marketplace with Sponsored Advertising," *Marketing Science*, 3 2024, *43* (2), 317–345.

Luca, Michael, "Reviews, Reputation, and Revenue: The Case of Yelp.Com," *SSRN Electronic Journal*, 3 2016.

_ **and Georgios Zervas**, "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud," *Marketing Science*, 1 2016, *62* (12), 3412–3427.

Marinescu, Ioana and Herbert Hovenkamp, "Anticompetitive Mergers in Labor Markets," *Indiana Law Journal*, 2019, *94* (3), 1031–1064.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier, "Promotional Reviews: An Empirical Investigation of Online Review Manipulation," *American Economic Review*, 8 2014, *104* (8), 2421–55.

McAfee, R Preston and Michael A Williams, "Horizontal Mergers and Antitrust Policy," *The Journal of Industrial Economics*, 1992, *40* (2), 181–187.

Miller, Nathan H and Matthew C Weinberg, "Understanding the Price Effects of the MillerCoors Joint Venture," *Econometrica*, 11 2017, *85* (6), 1763–1791.

Moraga-González, José Luis, Evgenia Motchenkova, and Saish Nevrekar, "Mergers and innovation portfolios," *The RAND Journal of Economics*, 12 2022, *53* (4), 641–677.

Motta, Massimo and Emanuele Tarantino, "The effect of horizontal mergers, when firms compete in prices and investments," *International Journal of Industrial Organization*, 9 2021, *78.*

_ **and Martin Peitz**, "Big tech mergers," *Information Economics and Policy*, 3 2021, *54*, 100868.

Neven, Damien J and Lars-Hendrik Röller, "Rent sharing in the European airline industry," *European Economic Review*, 1996, *40* (3), 933–940.

Nevo, Aviv, "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 2001, *69* (2), 307–342.

Newman, John M, "Antitrust in zero-price markets: Foundations," *University of Pennsylvania law review*, 2015, pp. 149–206.

Nocke, Volker and Nicolas Schutz, "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 3 2018, *86* (2), 523–557.

OECD, "DIRECTORATE FOR FINANCIAL AND ENTERPRISE AFFAIRS COMPETITION COMMITTEE Considering non-price effects in merger control-Background note by the Secretariat," Technical Report 2018.

_ , "Environmental Considerations in Competition Enforcement, Background Paper by the Secretariat," Technical Report, Paris 2021.

Ornaghi, Carmine, "Mergers and innovation in big pharma," *International Journal of Industrial Organization*, 2009, *27* (1), 70–79.

Parker, Geoffrey, Georgios Petropoulos, and Marshall Van Alstyne, "Platform mergers and antitrust," *Industrial and Corporate Change*, 10 2021, *30* (5), 1307–1336.

Perry, Martin K and Robert H Porter, "Oligopoly and the Incentive for Horizontal Merger," *The American Economic Review*, 1985, *75* (1), 219–227.

Prager, Elena and Matt Schmitt, "Employer Consolidation and Wages: Evidence from Hospitals," *American Economic Review*, 2 2021, *111* (2), 397–427.

Rochet, Jean Charles and Jean Tirole, "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 6 2003, *1* (4), 990–1029.

Schumpeter J A, *Capitalism, socialism and democracy* 1942.

Stiebale, Joel and Florian Szücs, "Mergers and market power: evidence from rivals' responses in European markets," *The RAND Journal of Economics*, 12 2022, *53* (4), 678–702.

Stigler Committee on Digital Platforms, "Final Report," Technical Report, Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business 2019.

Stigler, George J, Robert A Sherwin, Rebecca A Bissi, Anita Garten, William J Lynk, Lynn Shishido-Topel, and Robert S Stillman, "THE EXTENT OF THE MARKET," Technical Report 1985.

Sun, Monic, "How Does the Variance of Product Ratings Matter?," *Management Science*, 12 2011, *58* (4), 696–707.

Sweeting, Andrew, "The effects of mergers on product positioning: Evidence from the music radio industry," *RAND Journal of Economics*, 2010, *41* (2), 372–397.

Szücs, Florian, "M&A and R&D: Asymmetric Effects on acquirers and targets?," *Research Policy*, 2014, *43* (7), 1264–1273.

Tzanaki, Anna, "Dynamism and Politics in EU Merger Control: Appreciating the Gain and the Gap," *SSRN Electronic Journal*, 2023.

U.S. Department of Justice and Federal Trade Commission, "2023 Merger Guidelines," Technical Report, U.S. Department of Justice and Federal Trade Commission December 2023. Final joint guidelines issued 18 December

2023. Available at: `https://www.justice.gov/d9/2023-12/2023%20Merger%20Guidelines.pdf` (accessed 8 October 2025).

Valletti, Tommaso, "The innovation theory of harm in merger control: Some clarifications," *Economics Letters*, 9 2025, *255*, 112556.

Werden, Gregory J, "Demand elasticites in antitrust analysis," *Antitrust LJ*, 1997, *66*, 363.

Werden, Gregory J. and Luke M. Froeb, "The effects of mergers in differentiated products industries: Logit demand and merger policy," *Journal of Law, Economics, and Organization*, 10 1994, *10* (2), 407–426.

Williamson, O. E., "Economies as an Anti-Trust Defense: The Welfare Tradeoffs," in Charles K. Rowley, ed., *Readings in Industrial Economics: Volume Two: Private Enterprise and State Intervention*, London: Macmillan Education UK, 1972, pp. 111–135.

Yin, Pai-Ling, Jason P Davis, and Yulia Muzyrya, "Entrepreneurial Innovation: Killer Apps in the iPhone Ecosystem," *American Economic Review*, 2014, *104* (5), 255–259.

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