Essays on Product Portfolio Decisions and Welfare

by

Candice Wei Wang

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business and Economics)
in The University of Michigan
2026

Doctoral Committee:

Professor Ying Fan, Co-Chair
Professor Francine Lafontaine, Co-Chair
Professor Zach Brown
Professor Charles Murry

Candice Wei Wang
cweiwang@umich.edu
ORCID iD: 0009-0001-6350-6119

© Candice Wei Wang 2026

DEDICATION

To my parents, Chen Wang and Xiaoqiong Tan.

ACKNOWLEDGEMENTS

I would first like to express my deepest gratitude to the faculty who have supported me along the way, first and foremost my committee members: Ying Fan, Francine Lafontaine, Zach Brown, and Charlie Murry. I am profoundly grateful to Ying for her wisdom, patience, kindness, and the extraordinary advice, support, and guidance she has given me throughout this journey. Her course on industrial organization first sparked my interest in the field, and the rigor with which she approaches research has deeply shaped the way I think as an economist. I have learned so much from her sharp comments, her insights, and her remarkable ability to pinpoint the central issue with clarity and precision. I feel very fortunate to have learned from her. I am also extremely thankful to Francine for her emphasis on the bigger picture and her constant reminders to stay grounded in the real world. She has pushed me to think beyond the mechanics of the model, to consider what policymakers would care about, what the real-world analogues of my assumptions are, and how my methods and results should be interpreted in practice. This perspective has been invaluable, especially in structural IO, where complexity can sometimes make it easy to lose sight of economic and institutional realities. I am also grateful for her support and patience, and for the way she has always made space for me to think out loud in meetings, even when that meant going on tangents. I also thank Zach for his insightful advice, encouragement, patience, and careful feedback. I have especially appreciated his reminders to return to the data, look carefully at descriptive patterns, and use them to discipline the estimation, interpretation, and framing of my paper. His suggestions have often helped me step back from technical details and think more clearly about how to motivate the paper and communicate its main ideas. I am also very grateful for the detailed and constructive feedback he has given me after presentations, and for his patience and kindness whenever I ran into estimation difficulties, which happened often. Finally, I am grateful to Charlie for his advice, comments, and willingness to be direct. His questions pushed me to communicate my ideas more clearly, to be more explicit and careful about the assumptions and mechanisms in my models, and ultimately to develop a deeper understanding of my own work.

I am also extremely thankful to Joel Slemrod. Although my research eventually moved

toward industrial organization, Joel played an important role in my graduate school experience. The engaging and thought-provoking way he taught the field course drew me into tax research, and even after I moved away from tax, he continued to offer thoughtful comments on my research. His suggestions were often big-picture, but always sharp: he had an incredible intuition for identifying the key issues I needed to get right, as well as the parts I could let go. I also deeply appreciated his humor, which made meetings with him both fun and intellectually energizing. Beyond his academic guidance, I am especially grateful for his mentorship, and the kindness and support he has shown me over the years.

I would also like to thank the Business Economics program, the Department of Economics, and the Kresge Library for their generous grants and support. I am also grateful to the faculty in Economics and Business Economics for the many helpful comments and suggestions they provided during presentations and informal conversations.

I was incredibly fortunate to have met wonderful friends throughout this journey, who made graduate school so much more enjoyable. I am deeply grateful for the support they gave me through difficult times, for the many conversations we shared about research, life, and everything in between, and for the way we got through the hardest parts together. I will especially miss the lunches, the long conversations, and the Orlando trip with Niaoniao and Eva—what a ride that was. In particular, I would like to thank Aaron Kaye, Basil Isaac, Bei Luo, Bruna Guidetti, David Van Dijcke, Eric Parolin, Eva Wang, Gerardo Sanz-Maldonado, Harry Kleyer, Helen Wang, Jeongwon Jang, Jingyuan Cui, Joaquin Endara, Luciana Galeano, Meifeng Yang, Nathan Sunday, Niaoniao You, Peter Choi, Qize Chen, Rex Hsieh, Siyun He, Soyoung Han, Tammy Lee, Xiaomeng Li, Xinrui Zhou, Xinyu Chen, Yiman Ren, Yu Song, and Yuqing Gu. I would also like to thank my friend Lin Fan, who, although outside my program, has given me so much support throughout this journey. I am deeply grateful for the many conversations in which she helped me think through research ideas and gave me thoughtful feedback. I am equally grateful for her friendship, humor, and understanding, and for the way we got through difficult moments together. This journey was easier because of you.

Last but not least, I would like to thank my family for their immense support and understanding. I am extremely grateful to my mother, Xiaoqiong Tan, my father, Chen Wang, and my sister, Hannah Wang, for their love, encouragement, and support throughout this journey. Thank you for always believing in me, for giving me the freedom and confidence to pursue this path, and for being there for me through the ups and downs. Your support has been a constant source of strength, and I am beyond grateful.

I would also like to thank my husband, Jeff Zishen Cheng, for his love and unwavering support throughout this journey. I am deeply grateful for the countless rides to and from school, for the patience with which he listened to my late-night rants and half-formed research pitches, for the comfort he gave me through difficult moments, and for the way he encouraged me and celebrated the small victories along the way. His support has carried me through graduate school in more ways than I can count, and I genuinely could not have done this without him.

Finally, I would like to thank my cats, Creampuff and Barry, for their companionship and snuggles. Their quiet presence — though not always so quiet when they were hungry — made even the most stressful days feel a little softer.

TABLE OF CONTENTS

DEDICATION ii

ACKNOWLEDGEMENTS iii

LIST OF TABLES viii

LIST OF FIGURES x

LIST OF APPENDICES xi

ABSTRACT xii

CHAPTER

1 Complementarity and Product Variety: Strategic Product Offerings in Digital Ecosystems 1

1.1 Introduction 2
1.2 Industry Background and Data 7
1.2.1 Industry Background 7
1.2.2 Data 8
1.2.3 Quality Index 13
1.3 Model 16
1.3.1 Demand 16
1.3.2 Supply 17
1.4 Estimation 21
1.4.1 Demand and Marginal Cost 21
1.4.2 Fixed Costs 27
1.4.3 Estimation Results 29
1.5 Counterfactual Analysis 36
1.5.1 Counterfactual Results 40
1.6 Conclusion 46
References 48

2 Optimal Sin Tax with Endogenous Reformulation 50

2.1 Introduction 51

2.2 Theoretical Framework . . . 53
2.2.1 Individuals . . . 54
2.2.2 Firms . . . 55
2.2.3 Government . . . 56
2.2.4 Optimal Sin Tax Policy . . . 57
2.3 Empirical Model . . . 62
2.3.1 A Structural Approach . . . 62
2.3.2 Demand . . . 63
2.3.3 Supply . . . 64
2.4 Data . . . 65
2.5 Estimation . . . 68
2.6 Estimation Results . . . 69
2.7 Reformulation Cost Calibration . . . 74
2.8 Optimal Linear Sugar Taxes . . . 77
2.9 Conclusion . . . 82
References . . . 83

APPENDICES . . . 86

LIST OF TABLES

TABLE

1.1 IDC Product-Market Summary Statistics . . . 11
1.2 Complementarity Features and Rollout Timing . . . 12
1.3 PCA Loadings for the Quality Index . . . 13
1.4 Monte Carlo Validation of Ownership-State Simulation . . . 23
1.5 Choice-Set Variation Along the Quality Index . . . 26
1.6 Demand Estimates . . . 30
1.7 Selected Micro-Moment Fit . . . 31
1.8 Demand Diversion Ratios . . . 34
1.9 Marginal-Cost Estimates . . . 34
1.10 Fixed-Cost Parameter Estimates . . . 35
1.11 Counterfactual Changes in Product Variety, Prices, and Welfare . . . 40
1.12 Welfare Decomposition by Counterfactual Channel . . . 41
1.13 Product-Portfolio Response Under No Complementarity . . . 42
1.14 Product-Portfolio Response Under Existing-Feature Interoperability . . . 43
1.15 Product-Portfolio Response Under Full-Feature Interoperability . . . 45

2.1 Toy Mechanism Exercise: Welfare Comparison of Linear and Notched Taxes . . . 62
2.2 Average Monthly Consumption, Trip, and Retailer Statistics . . . 66
2.3 Product, Brand, Market-Share, and Volume Statistics by Market . . . 67
2.4 Product, Brand, Market-Share, and Price Statistics by Product Category . . . 67
2.5 Demand Estimates . . . 69
2.6 Pooled Grouped Own- and Cross-Price Elasticities . . . 71
2.7 Marginal Cost Estimates . . . 72
2.8 Micro Moments: Observed vs. Model-Implied . . . 73
2.9 Reformulation Cost Calibration . . . 76
2.10 Optimal Tax Choice: Fixed Characteristics vs. Endogenous Reformulation . . . 78
2.11 Tax Choice When Reformulation Is Available . . . 79

A.1.1 Statista Consumer Insights Sample for Ownership Moments . . . 86
A.1.2 Full Micro-Moment Fit . . . 87
A.1.3 OpenBrand Survey Sample for Income Moments . . . 89
A.1.4 Highest-Quality Smartphone and Tablet Products, 2022Q4 . . . 90

A.1.5 Demand Responses to a $10 Price Increase 91
A.4.1 Monte Carlo Recovery of Standard Demand Parameters 101

B.1.1 Toy Simulation Primitives 108

LIST OF FIGURES

FIGURE

1.1 Number of Unique Models and Feature Rollouts by Vendor 9
1.2 Excess Joint Ownership Relative to Independence, 2022 9
1.3 Evolution of the Quality Index 14
1.4 Smartphone Quality-Space Variety by Firm Group 15
1.5 Tablet Quality-Space Variety by Firm Group 15
1.6 Inside-Good Taste Heterogeneity and Ownership Stocks 27
1.7 Product-Level Fixed-Cost Bounds: Smartphones 37
1.8 Product-Level Fixed-Cost Bounds: Tablets 37

2.1 Sugar Content of In-Sample Products, 2018 68
2.2 Distribution of implied markups 74
2.3 Distribution of implied marginal costs 74
2.4 Welfare Decomposition for Linear Sugar Taxes 80
2.5 Sugar Content Under the Fixed-Characteristics Optimum 81
2.6 Sugar Content Under the Endogenous-Reformulation Optimum 81

A.1.1 Ownership Rates Used in the Ownership Micro-Moments 92
A.1.2 Minimum-to-Maximum Quality-Index Range by Brand 93
A.1.3 Range of Screen Sizes by Brand 93
A.1.4 Range of Pixel Density by Brand 94
A.1.5 Range of RAM by Brand 94
A.1.6 Range of Processor Speed by Brand 95
A.1.7 Distribution of Recovered Marginal Costs 96
A.5.1 Example Potential-Product Quality Draws 105

LIST OF APPENDICES

APPENDIX
Appendix A. Complementarity and Product Variety: Appendix 86
Appendix B. Optimal Sin Tax with Endogenous Reformulation: Appendix 108

ABSTRACT

This dissertation consists of two chapters in Industrial Organization that study firm choices beyond prices. The chapters examine product portfolios in digital ecosystems and reformulation incentives under corrective taxation. Together, they show how demand linkages and regulation can shape firms' product offerings and product attributes.

Chapter 1, "Complementarity and Product Variety: Strategic Product Offerings in Digital Ecosystems," studies how ecosystem complementarities affect firms' product choices and the welfare consequences of interoperability policies. I develop and estimate a structural model of demand, pricing, and product choice in the U.S. smartphone and tablet markets from 2010 to 2022. The estimates imply willingness to pay of about $47 for a smartphone compatible with a consumer's owned tablet and $51 for a tablet compatible with a consumer's owned smartphone. Relative to the observed equilibrium, removing demand-side cross-device complementarity affects product variety through portfolio reallocation rather than through a simple change in aggregate product counts: firms with established ecosystems reduce their product portfolios, while rivals expand into the resulting product-space gaps. This reallocation lowers quality-space variety by 3.0 percent and total welfare by 17.8 percent. A government mandate requiring interoperability across existing features generates welfare gains through broader compatibility and lower prices, but these gains are partially offset by induced portfolio responses, as ecosystem incumbents reduce their portfolios and rival firms add products that do not offer the full set of complementarity features. In contrast, under a full-feature interoperability benchmark, in which all complementarity features are available on all models and interoperable across device pairs, rival product entry preserves complementarity value and the portfolio response becomes welfare improving. In the 2022Q3 market, relative to the observed equilibrium, existing-feature interoperability raises welfare by $0.3 billion, with portfolio reoptimization offsetting about 70 percent of direct compatibility and price-response gains; full-feature interoperability raises welfare by $4.3 billion.

Chapter 2, "Optimal Sin Tax with Endogenous Reformulation," studies corrective taxation when firms can redesign the taxed product attribute. The chapter develops an optimal

tax framework in which firms choose both prices and sugar content. Using data from the California non-alcoholic beverage market, we estimate a structural model of demand and supply and calibrate reformulation costs using evidence from the U.K. Soft Drinks Industry Levy. Under an efficiency benchmark that places equal welfare weight on all consumers and therefore abstracts from redistributional concerns, incorporating manufacturers' reformulation responses nearly triples the optimal linear sugar tax on a 12 g/100ml drink, increases annual welfare gains from $39.9 million to $261.6 million, and lowers average purchased sugar from 6.27 to 1.74 g/100ml. The results show that analyses treating product characteristics as fixed can understate both the gains from sugar taxation and the optimal tax rate.

CHAPTER 1

Complementarity and Product Variety: Strategic Product Offerings in Digital Ecosystems

Abstract. Complementarity across product markets can shape firms' strategic product offerings by generating demand spillovers across compatible products. A marginal product can generate value not only through its own sales, but also by expanding the installed base that raises demand for complementary products in adjacent markets. Interoperability alters firms' ability to appropriate these cross-market spillovers, so its welfare effects depend not only on prices and consumers' direct gains from compatibility, but also on firms' endogenous product-portfolio responses. This paper studies this mechanism in U.S. smartphone and tablet markets from 2010 to 2022. I develop and estimate a structural model of demand, pricing, and product choice in which consumers value cross-device features when owned devices are compatible and firms choose product portfolios internalizing own-category profits and follow-on demand in the adjacent category. The estimates imply willingness to pay of about $47 for a smartphone compatible with an owned tablet and $51 for a tablet compatible with an owned smartphone. Counterfactuals show that complementarity affects product portfolios mainly through reallocation across firms. Removing complementarity lowers welfare by $1.6 billion, reduces aggregate product count by 1.3 products, and lowers quality-space variety by 38.2 units: Apple and Samsung reduce their product offerings, while other firms expand into the resulting product-space gaps. Making existing features interoperable increases product count by 3.3 products, raises quality-space variety by 22.7 units, and raises welfare by $0.3 billion; however, portfolio reoptimization reduces welfare by $0.67 billion, offsetting about 70 percent of the direct compatibility and price-response

gains. A broader full-feature interoperability benchmark raises welfare by $4.3 billion because rival entry then preserves complementarity value rather than merely replacing products that exit.

1.1 Introduction

Complementarity arises when the utility one derives from a product depends on the ownership or use of another. While some complementarity naturally exists across all products within their respective categories, such as milk and cereal, others appear only for certain product pairs. A prominent setting in which the latter occurs is digital ecosystems: groups of interconnected devices, software, and services linked by firm-developed features and compatibility rules. For example, features such as cloud synchronization, wireless sharing, and continuity raise the value of owning smartphones and tablets within the same ecosystem by enabling consumers to move content and tasks seamlessly across products.

When complementarity arises only between compatible product pairs offered by the same firm, it changes firms' return to product variety. In the absence of complementarity, firms determine product variety by weighing the incremental profit from the additional product against its fixed cost, net of cannibalization. However, when products are interconnected through complementarity, introducing a new product can raise profit not only through its own sales, but also by increasing demand for compatible products offered by the same firm in adjacent categories. These cross-product demand spillovers, particularly in the presence of heterogeneous consumer preferences, can substantially alter firms' variety decisions by amplifying the returns to carrying a more differentiated portfolio. Apple's introduction of the iPhone SE illustrates this margin: in its 2016Q2 earnings call, Apple described the model as attracting more "new-to-iPhone customers" and bringing "new people into the iOS ecosystem."

This product-variety channel also changes the welfare analysis of digital ecosystems and interoperability policies. Policy concerns often emphasize the market-power effects of limited compatibility: ecosystem-specific features can reduce substitution across firms, increase switching costs, and allow dominant firms to raise prices. Recent policy interventions reflect this concern, including the DOJ's 2024 Apple suit and the EU Digital Markets Act's interoperability obligations for designated gatekeepers.[1] However, when product portfolios

[1] In 2024, the U.S. Department of Justice and state attorneys general sued Apple, alleging that restrictions around the iPhone ecosystem impede switching and interoperability. In the European Union, the Digital

are endogenous, interoperability also changes firms' incentives to introduce and maintain products.[2] On the one hand, by making cross-device features work across ecosystems, interoperability reduces the extent to which demand spillovers are tied to a firm's own ecosystem. A firm may therefore have weaker incentives to offer products whose profitability depends on generating follow-on demand within its own ecosystem. On the other hand, interoperability can raise the incremental return to differentiated products by making them attractive to consumers outside the firm's existing ecosystem. For example, a Samsung tablet may become more appealing to iPhone owners once cross-device features work across ecosystems, strengthening incentives to introduce products that reach otherwise underserved consumer segments. The net effect on product variety depends on whether the added demand from rival-ecosystem owners exceeds the lost value of within-ecosystem follow-on demand, after accounting for cannibalization, price competition, and fixed costs. The welfare effect of interoperability therefore depends not only on direct consumer gains from broader compatibility and lower prices, but also on how firms reoptimize their product portfolios when intra-ecosystem spillovers are replaced by broader inter-ecosystem spillovers.

This paper studies how ecosystem complementarity affects firms' product choices and how endogenous product variety changes the welfare consequences of interoperability policy. I focus on three questions. First, to what extent does ecosystem complementarity increase the number and differentiation of products firms offer? Second, how does complementarity affect equilibrium prices and consumer welfare once product portfolios adjust? Third, how do interoperability mandates affect welfare when firms reoptimize both prices and product offerings?

I answer these research questions by developing a structural model of demand, pricing, and product choice in the context of the U.S. smartphone and tablet markets. On the demand side, consumers choose smartphones and tablets and derive utility from product characteristics, prices, replacement incentives, and complementarity with their owned devices. On the supply side, firms first choose product portfolios to maximize expected profits, internalizing contemporaneous profits within the product's own category and follow-on profits generated by complementarity-driven demand spillovers to the adjacent category. After product offerings are chosen, firms observe demand and cost shocks and set prices.

Markets Act requires designated gatekeepers to provide interoperability with operating-system hardware and software features, and the European Commission adopted specification decisions applying these obligations to Apple devices in 2025.

[2]In this paper, I use interoperability to refer to settings or policies in which cross-device features previously available or compatible only within the same ecosystem can operate across ecosystems, allowing consumers to realize complementarity benefits from device pairs supplied by different firms.

I estimate the demand and pricing model jointly using U.S. smartphone and tablet products sold between 2010 and 2022. The main dataset comes from International Data Corporation (IDC)'s Mobile Phone Tracker and PC Tracker, which report quarterly prices, quantities, and hardware characteristics for products sold in the United States. To measure hardware quality, I construct category-specific quality indices using the first principal component of standardized product attributes separately within smartphones and tablets. I supplement these product data with manually collected information on cloud synchronization, wireless sharing, and continuity features. Finally, I use Statista and Pew measures of overall device ownership and cross-device joint-brand ownership, together with OpenBrand purchase surveys, to construct micro-moments that inform income heterogeneity in demand and discipline device ownership patterns. Demand estimates show that consumers on average value device quality, with especially large heterogeneity in preference for tablet quality. Consumers also value same-ecosystem compatibility: evaluated at median income, the estimated complementarity index implies willingness to pay of about $47 for an additional smartphone when the consumer owns a compatible tablet and about $51 for an additional tablet when the consumer owns a compatible smartphone. On the supply side, recovered marginal costs rise with hardware quality and model age.

Using the no-profitable-deviation conditions implied by Nash equilibrium in product offerings, I recover lower bounds on fixed costs of offered products and upper bounds on fixed costs of not-offered feasible products. I then parameterize fixed costs as functions of product quality, product age, brand, and distance from the firm's nearest existing product in quality. Following the bounded-support logic in Eizenberg (2014), I complete the one-sided bounds using comparable products within cells defined by product and portfolio attributes, and use the resulting intervals to set-identify fixed-cost parameters. The estimates imply high fixed-cost components for Apple and Samsung products, and fixed costs declining as products age.

Using the model estimates, I conduct counterfactuals that vary the complementarity environment linking smartphones and tablets. The first counterfactual removes Cloud Sync, Share, and Continuity for all products, isolating the role of ecosystem complementarity in the observed equilibrium. Removing complementarity lowers total welfare by $1.6 billion in the 2022Q3 market. Most of this loss comes from the direct value of cross-device features, although lower prices partially offset it. Portfolio responses are more mixed; while firms with established ecosystems reduce product offerings, rivals expand into the resulting gaps in product space, resulting in a modest change in aggregate product count.

The second counterfactual evaluates existing-feature interoperability: active cross-device features are made interoperable across ecosystems, while each product's observed complementarity feature availability is held fixed. Interoperability has opposing effects on firms' product-choice incentives. On the one hand, it weakens within-ecosystem follow-on demand, reducing the value of marginal products for firms with established ecosystems. On the other hand, it makes rival products more attractive to consumers who already own devices in another ecosystem, increasing the potential return to differentiated products offered by other firms. In the estimated 2022Q3 market, relative to the observed equilibrium, the second force increases product offerings by Samsung and other non-Apple firms, raising aggregate product counts and quality-space variety. For Apple, however, the first force dominates, leading it to narrow its product offerings. Relative to the observed 2022Q3 equilibrium, existing-feature interoperability raises welfare by \$0.3 billion, with gains coming from broader compatibility and lower prices. These gains are partially offset by the endogenous portfolio response: despite the increase in aggregate product counts and quality-space variety, the products added or retained under interoperability do not fully compensate consumers and firms for the products that exit, especially in smartphones.

I therefore also consider a broader full-feature interoperability benchmark in which all products offer all cross-device complementarity features and these features are fully interoperable across ecosystems. This benchmark raises welfare by \$4.3 billion and generates a positive portfolio channel. Comparing the two interoperability exercises clarifies the mechanisms behind the ambiguous supply-side effects. When feature availability is held fixed, interoperability primarily reallocates demand across products that already differ in their feature environments; product variety can rise, but the portfolio response can still reduce welfare if the entering products do not replace the value of the products that exit. When expanding compatibility is accompanied by broader adoption of cross-device features, rival entry becomes more valuable because added products also carry complementarity benefits. In that case, the portfolio response becomes a major source of welfare gains.

This paper contributes to the literature on indirect network effects and complementarities across markets. Early empirical work studies settings in which one side has limited standalone value without the complementary good, such as CD players and titles (Gandal, Kende, and Rob 2000), video game consoles and exclusive titles (Lee 2013), and electric vehicles and charging stations (Li 2019). Other work studies products with substantial standalone value whose utilities nevertheless depend on joint consumption, including online and print newspapers (Gentzkow 2007) and asymmetric product complements (Lee, Kim,

and Allenby 2013). Smartphones and tablets fit the latter class. Each product is valuable on its own, but cloud synchronization, wireless sharing, and continuity features make ownership of one device affect demand for the other. The setting adds two margins that are central for policy: the goods are durable, so past purchases create ownership stocks that determine current complementarity benefits; and the leading firms operate in both markets, so a marginal product can generate profits through future demand for compatible products in the other category.

Recent work on mobile ecosystems shows that these demand-side forces are economically large. Lidany (2025) estimates smartphone–laptop compatibility and compares open and closed ecosystems; his experimental evidence implies that compatibility raises willingness to pay for a compatible smartphone by about \$75, or 9% of the retail price. Chun (2026) estimates a dynamic model of mobile bundles and evaluates remedies in the DOJ case against Apple; he estimates the cost of switching from iOS to Android at about \$93, or 45.8% of the mean purchase price. These papers show that ecosystem attachment, compatibility, and switching frictions can meaningfully shape demand and pricing. My paper studies the supply-side margin that remains fixed in those analyses: when compatibility affects the value of an installed base, firms may change which products they introduce and maintain.

The paper also contributes to the literature on endogenous product choice. Fan (2013), Eizenberg (2014), Fan and Yang (2020), and Wollmann (2018) show that policy counterfactuals can change substantially once firms adjust product portfolios. Fan and Yang (2020) apply this logic to smartphone mergers, where product entry and exit mediate the welfare effects of ownership changes. In my setting, the product-choice margin operates across two durable-goods markets. A phone can affect future tablet demand through the installed base of compatible phone owners, and a tablet can affect future phone demand in the same way. Interoperability therefore changes firms' ability to appropriate cross-market spillovers, so its welfare effects depend on both the direct compatibility gains and the induced response in product variety.

The closest paper in this strand is Chatterjee et al. (2024), which studies cross-market spillovers and endogenous product portfolios in complementary markets. They show that early entry of technologically advanced 4G handsets speeds up wireless network rollout, and that faster network rollout raises the profitability and variety of 4G handsets. Consumers benefit from diffusion in both markets, while domestic handset firms face a tradeoff between direct competition from international firms and the indirect spillover created by a more developed complementary market. This paper shares the idea that complementar-

ity can raise product variety, but the mechanism is different. In Chatterjee et al. (2024), product variety is tied to technology diffusion: more advanced-generation handsets become profitable as complementary network coverage expands, and the complementarity operates through same-period handset-service availability. Here, the product-variety margin comes from durable-goods lock-in and heterogeneous preferences. A differentiated phone or tablet can attract a consumer segment into the ecosystem, and that installed base then raises future demand for compatible devices. Complementarity and durability therefore magnify the return to variety: firms value marginal products not only for their own sales, but also for the future compatible purchases of consumers they bring into the ecosystem. This makes interoperability the central counterfactual margin: changing compatibility across ecosystem boundaries changes consumers' direct compatibility benefits and firms' ability to appropriate the future demand spillovers that support marginal products.

The remainder of the chapter proceeds as follows. Section 1.2 provides background on the U.S. smartphone and tablet markets and describes the data. Section 1.3 presents the demand, pricing, and product-choice model. Section 1.4 describes estimation and reports the demand, marginal-cost, and fixed-cost estimates. Section 1.5 evaluates no-complementarity and interoperability counterfactuals. Section 1.6 concludes.

1.2 Industry Background and Data

1.2.1 Industry Background

This paper studies U.S. smartphone and tablet markets from 2010 to 2022, a period that captures the early growth and subsequent maturation of both industries. Smartphones are one of the largest consumer-electronics categories, with over 120 million units sold in 2022 and total sales value of $101.7 billion. Tablets are smaller but still substantial, with over 26 million units sold in 2022 and revenue of $15.7 billion.

Both product categories emerged in the late 2000s. Early smartphones were produced by firms such as Nokia and BlackBerry, but Apple's release of the first iPhone in 2007 is widely regarded as the beginning of the modern smartphone era. The modern tablet market began with Apple's launch of the iPad in 2010, which quickly spurred industry-wide adoption.

The two markets exhibit substantial firm overlap, creating scope for cross-market strategy. Of the 47 firms active in at least one market, 19 operated in both smartphones and tablets. These dual-market firms account for 98% of smartphone sales and 92% of tablet

sales in the dataset. The industries are also highly concentrated, with Apple and Samsung jointly capturing the majority of sales in both categories throughout the sample period. To compete across both markets, these firms developed ecosystem-specific complementarity features, including cross-device file sharing, app continuity, and cloud-based syncing. Apple integrated these features through its closed operating system, while Samsung layered them onto Android through pre-installed applications. Both firms, and especially Apple, also expanded product portfolios to reach broader consumer segments and leverage the complementarities generated by these features. Apple introduced lower-cost models such as the iPhone SE and diversified its tablet offerings across the iPad Mini, Air, and Pro lines.

Figure 1.1 shows that the leading firms in each category have expanded their product offerings over the sample, while industry-median counts remain low. The dashed vertical lines mark the staggered rollout of the ecosystem features in Table 1.2, making the timing of product-line expansion visible relative to the introduction of cross-device compatibility. Apple and Samsung carry many more concurrent smartphone and tablet models than the typical competitor, consistent with strategies that rely on broad portfolios across ecosystem segments, motivating the central question studied in this paper.[3]

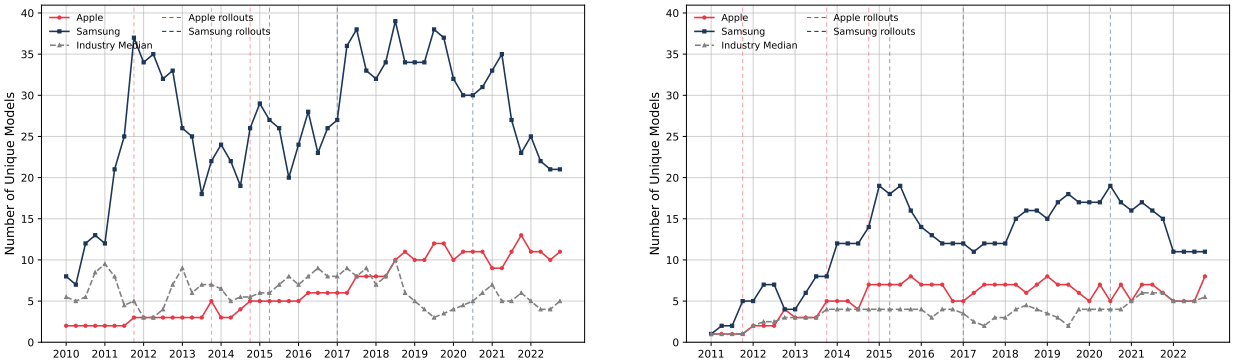
Joint-ownership patterns provide a complementary motivation for studying ecosystem complementarity. Figure 1.2 reports the difference between observed joint ownership shares and the product of marginal shares, $P(\text{phone}, \text{tablet}) - P(\text{phone})P(\text{tablet})$, using 2022 Statista joint-ownership data. Positive cells along the diagonal for Apple and Samsung indicate that within-ecosystem joint ownership is substantially more common than independent product choice would imply, while off-diagonal cells are typically negative. This co-ownership pattern is consistent with both correlated tastes and ecosystem complementarity; the structural model uses staggered feature rollouts to separate these mechanisms.

1.2.2 Data

I combine several data sources to measure consumer demand, ownership patterns, complementarity and compatibility.

Product-market data.— The primary data come from the International Data Corporation (IDC) Mobile Phone Tracker and PC Tracker, which report prices, units sold, and hardware

[3]Samsung adopted a many-product strategy early on and offers substantially more products than Apple. This difference is consistent with Samsung's greater vertical integration in manufacturing: Samsung operates its own production facilities and component businesses, thereby reducing the incremental fixed costs of introducing and offering additional products, whereas Apple relies primarily on external suppliers for manufacturing (Samsung Electronics Co., Ltd. 2024, pp. 4, 18; Apple Inc. 2024, pp. 5, 7, 46).



(a) Smartphones

(b) Tablets

Figure 1.1: Number of Unique Models and Feature Rollouts by Vendor

Notes: The figure plots the number of distinct models offered each quarter by Apple, Samsung, and the industry median across firms. Models are counted at the model-quarter level used in estimation. Dashed vertical lines mark the initial rollout quarters for the complementarity features in Table 1.2: red for Apple (2011Q4, 2013Q4, 2014Q4) and blue for Samsung (2015Q2, 2017Q1, 2020Q3). The tablet panel uses the same y-axis range as the smartphone panel to facilitate cross-market comparison.

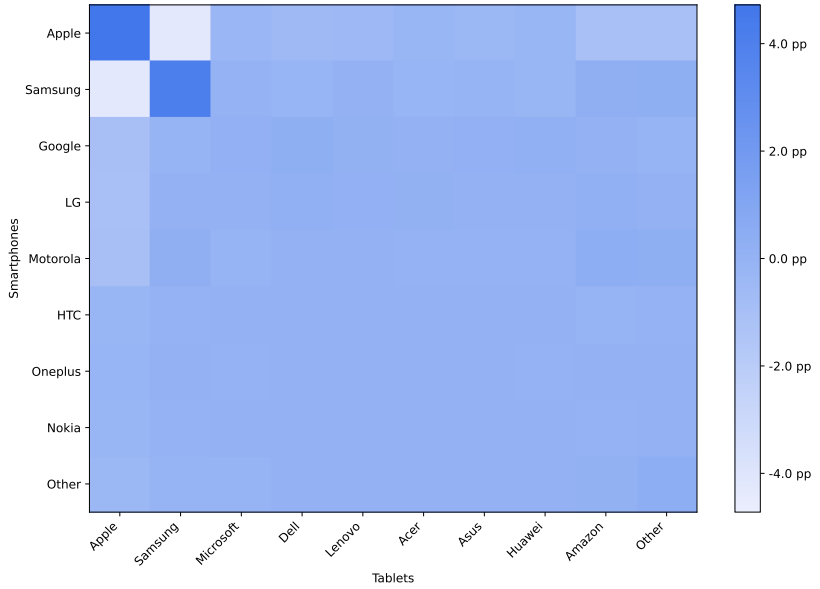


Figure 1.2: Excess Joint Ownership Relative to Independence, 2022

Notes: Cells report $P(\text{phone brand}, \text{tablet brand}) - P(\text{phone brand})P(\text{tablet brand})$ in percentage points, using 2022 cross-brand ownership shares from the Statista U.S. Consumer Survey. Marginals are computed across the displayed brand cells. Positive values indicate that the joint ownership combination is more common than independence implies; negative values indicate the opposite. This cross-sectional pattern is descriptive; the staggered rollout timing in Table 1.2 and Figure 1.1 provides time-series variation used by the model to distinguish complementarity from correlated brand tastes.

characteristics for smartphones and tablet PCs sold in the United States.[4] The final demand sample covers 2010Q1–2022Q4 for smartphones and 2011Q1–2022Q4 for tablets.[5] Products are defined at the model level. For example, iPhone 4S storage configurations are combined into a single iPhone 4S product in each quarter.[6]

Table 1.1 summarizes the final IDC product-market samples used in estimation, including product counts, prices, quantities, the quality index, and the hardware characteristics used to construct it.

Complementarity data.— I supplement the IDC data with manually collected public information on complementarity features and compatibility across devices. I focus on three feature families: Cloud Sync, Share, and Continuity. Table 1.2 summarizes the function, availability, and initial rollout timing in the estimation sample. The release timing is staggered across feature families and ecosystems: Apple introduces Cloud Sync and Share before comparable Samsung or Android services, while Continuity appears in adjacent years for Apple and Samsung. I measure feature availability at the product-quarter level; availability and compatibility can begin and change after a device's release upon software support update. A smartphone–tablet pair receives complementarity from a feature only when public support information indicates that both devices can use the feature with each other in that quarter.[7] Complementarity therefore varies across firms over time, and across the specific device pairs owned.

Cost shifters and market sizes.— To construct cost shifters and market sizes, I use Bureau of Labor Statistics data on population, import price indices, and exchange rates. I also use historical memory and flash memory prices from an online dataset.

[4]I supplement missing hardware characteristics using GSMArena and additional web-scraped or manually collected sources.

[5]The demand sample includes smartphone observations and consumer-segment tablet observations. I exclude carrier-branded phones since their prices are often tied to service-contract incentives not observed in my data. I exclude nonstandard tablet products, such as kids' tablets, keyboard-bundled tablets, and Windows or ChromeOS tablets, as they serve distinct use cases from slate tablets such as the iPad and Samsung Galaxy Tab, which form the core tablet segment through which the ecosystem complementarity features operate. I also drop tablet records reported under the aggregate brand category "Others," product-quarter observations with fewer than 100 units sold, and brands whose maximum quarterly inside share is below 0.1 percent; together, these exclusions account for less than 2 percent of total sales.

[6]Models are defined as brand–model name–hardware-feature combinations, with storage and RAM configurations of the same device model aggregated into a single product. Storage and RAM are retained as unweighted average product characteristics to avoid treating demand-driven changes in the mix of configurations sold as changes in the product, while price is measured as the sales-weighted average across SKU-quarters.

[7]I collect rollout dates and device-support information from public feature-release announcements, product documentation, and device-support pages.

Table 1.1: IDC Product-Market Summary Statistics

Variable	N	Mean	SD	Min	Median	Max
Panel A. Smartphones						
Products per quarter	52	192.1	82.9	64.0	193.0	349.0
Units sold (000s)	9,987	182.5	674.1	0.1	26.6	14,874.0
Real price	9,987	253.1	225.5	9.9	165.8	2,059.4
Quality index	9,987	0.00	2.53	-7.58	-0.16	6.28
Screen size	9,987	5.08	1.02	2.20	5.00	8.30
OLED display	9,987	0.197	0.398	0.000	0.000	1.000
Pixel density	9,987	3.1	1.0	1.2	2.8	8.0
Processor speed	9,987	1.59	0.60	0.30	1.40	3.50
Processor cores	9,987	4.69	2.53	1.00	4.00	10.00
Storage (GB)	9,987	42.7	74.2	0.1	16.0	768.0
RAM (GB)	9,987	2.39	2.33	0.02	2.00	18.00
Network generation	9,987	1.80	0.56	0.00	2.00	3.00
Rear cameras	9,987	1.35	0.76	0.00	1.00	5.00
Thickness	9,987	9.62	2.09	4.80	9.10	23.00
Panel B. Tablets						
Products per quarter	48	80.5	24.9	12.0	78.0	126.0
Units sold (000s)	3,863	118.8	487.7	0.1	8.2	11,400.0
Real price	3,863	202.1	170.2	25.7	141.1	1,303.1
Quality index	3,863	-0.00	2.22	-4.28	-0.40	8.66
Screen size	3,863	8.55	1.47	7.00	8.00	14.60
OLED display	3,863	0.027	0.162	0.000	0.000	1.000
Pixel density	3,863	1.8	0.6	0.9	1.7	3.6
Processor speed	3,863	1.47	0.45	0.60	1.30	3.49
Processor cores	3,863	3.53	2.15	1.00	4.00	8.00
Storage (GB)	3,863	30.4	70.4	1.0	16.0	793.6
RAM (GB)	3,863	1.54	1.36	0.10	1.00	12.00
Rear cameras	3,863	0.92	0.33	0.00	1.00	2.00
Thickness	3,863	9.69	2.05	4.70	9.90	20.00

Notes: The table reports summary statistics for the final IDC product-market samples used in estimation. The smartphone sample contains 9,987 product-quarter observations in 52 quarters, and the tablet sample contains 3,863 product-quarter observations in 48 quarters. Products per quarter is computed at the quarterly product-market level; all other rows are product-quarter observations. Units are reported in thousands. Prices are real prices used in estimation. Storage and RAM are reported in GB; tablet storage is converted back from the scaled estimation variable for readability. The variable set includes unit sales, the PCA quality index, and the hardware characteristics used to construct the quality index.

Ownership moments.— I use Statista Consumer Insights and Pew Research Center surveys to discipline ownership states. Statista Consumer Insights is a repeated online survey of U.S. internet users aged 18–64, fielded in rolling waves with quotas on age, gender, and region. The Statista extracts report aggregate tables rather than individual panel histories: smartphone users report the brand of their primarily used smartphone, tablet-owning households report the tablet brands they own, and cross-tabulations report tablet brand ownership by smartphone brand. I use the 2019–2022 Statista brand and cross-brand tables to construct brand-ownership and joint-ownership moments. The final Statista moments

Table 1.2: Complementarity Features and Rollout Timing

Feature	Description	Ecosystem	Initial rollout
Cloud Sync	Syncs information across devices.	Apple, Samsung	Apple: 2011Q4 Samsung: 2017Q1
Share	Sends files, photos, or links across devices.	Apple, Samsung, Android	Apple: 2013Q4 Samsung/Android: 2020Q2–2020Q3
Continuity	Lets users continue tasks, calls, texts, or application sessions across devices.	Apple, Samsung	Apple: 2014Q4 Samsung: 2015Q2

Notes: Rollout timing is measured by the first quarter in which supported phone-tablet pairs appear in the estimation sample. Dates come from public feature-release announcements and device-support information. A feature is available for a phone-tablet pair only when public support information indicates that both products support the same feature in that quarter.

use 41,780 brand-ownership observations and 17,565 cross-ownership observations.[8] Pew Research Center surveys provide annual smartphone and tablet ownership rates that anchor population-wide device ownership levels.

Income moments.— I use OpenBrand's Consumer Electronics survey to discipline income heterogeneity in demand. OpenBrand surveys U.S. households quarterly about durable-goods purchases made in the previous 90 days, including the product purchased, brand, transaction price, retailer, alternatives considered, and stated reasons for purchase. I use the 2015–2022 records for cell phones and tablets, which report survey weights and respondent income bins.[9] I aggregate these records into income-specific moments that summarize differences in brand choice and paid prices across consumers, as described in Section 1.4.1.2. These moments complement the aggregate IDC market shares and ownership moments by providing additional information on income heterogeneity in price sensitivity and ecosystem preferences. Appendix Table A.1.3 reports summary statistics for the OpenBrand sample used in the income moments.

[8] Appendix Table A.1.1 reports the underlying sample sizes by survey wave and the pooled brand composition.

[9] The income-moment sample keeps priced cell-phone and tablet purchases with nonmissing brand, income bin, and transaction price. I map reported income ranges into Low, Middle, and High bins and deflate paid prices by the quarterly CPI index normalized to 2010Q1=100 before constructing high-price indicators.

1.2.3 Quality Index

In order to capture heterogeneity in preferences for product quality without estimating random coefficients on a large vector of hardware characteristics, I summarize observed hardware characteristics using category-specific quality indices. Specifically, within each product category, I standardize the attributes listed in Table 1.3 and use their first principal component as a summary measure of vertical quality. This approach captures common hardware upgrading without requiring separate coefficients on many highly correlated characteristics in a setting with limited market-level variation. The resulting index provides a parsimonious measure of hardware quality for demand, marginal cost, and product-choice payoffs.[10] Table 1.3 reports the PCA loadings, and Figure 1.3 shows broad upgrading in smartphones and tablets over the sample.

Appendix Table A.1.4 makes the scale of the index concrete by listing the ten highest-quality smartphones and tablets in 2022Q4, together with selected hardware attributes.

Table 1.3: PCA Loadings for the Quality Index

Attribute	Smartphone loading	Tablet loading
Screen size	0.351	0.262
OLED display	0.193	0.191
Pixel density	0.304	0.354
Processor speed	0.359	0.387
Processor cores	0.335	0.369
Storage	0.365	0.381
RAM	0.376	0.408
Network generation	0.312	
Rear camera	0.262	0.219
Thickness	-0.258	-0.352
Share of variance explained	0.642	0.546

Notes: The table reports the loadings from the first principal component, computed separately for smartphones and tablets. Prior to constructing the index, each characteristic is standardized within product category to have mean zero and standard deviation one. The reported loadings are the coefficients on these standardized characteristics in the first principal component. The negative loading on thickness is a feature of the estimated first component: thinner devices tend to vary with higher values of the other quality-related hardware attributes.

[10]PCA loadings reflect the covariance structure of observed hardware attributes across products. An alternative approach by Fan (2013) and Fan and Yang (2020) constructs quality indices as linear combinations of product characteristics with weights estimated in the demand model; the relative importance of each characteristic is then determined through observed sales patterns. In my setting, however, this approach would add a large number of nonlinear parameters to the estimation, substantially expanding the search space and making the procedure computationally challenging.

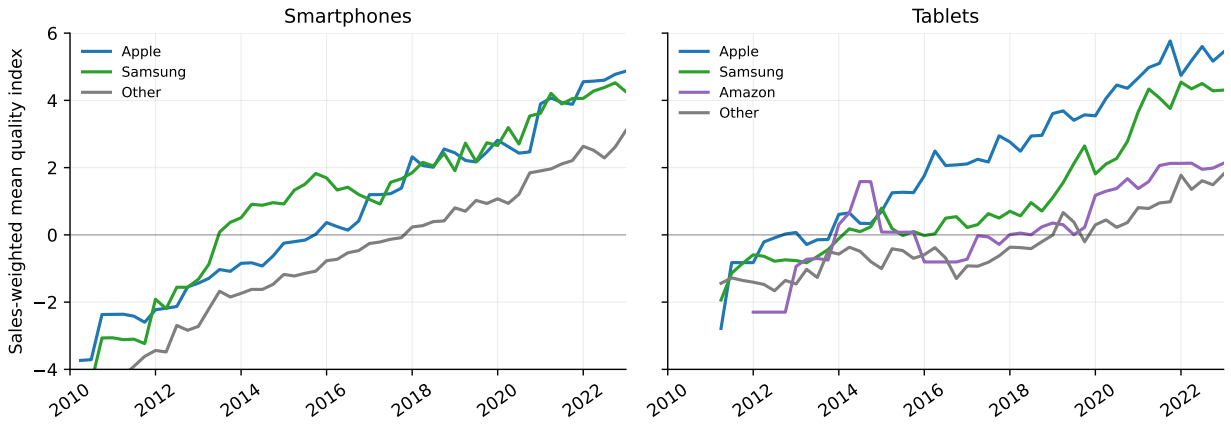


Figure 1.3: Evolution of the Quality Index

Notes: The figure plots quarterly sales-weighted mean values of the quality index by product category and major brand group. The index is constructed separately for smartphones and tablets from standardized hardware characteristics.

Within each quarter, leading vendors span a significant portion of the industry-wide quality space. I summarize this margin using the Fan–Yang quality-spacing measure in Fan and Yang (2020). For each firm-category-quarter, products are sorted by the quality index; the measure sums the square roots of adjacent positive quality gaps and squares the sum, assigning higher variety to portfolios that cover a wider range of quality levels and fill that range with more distinct products. Figures 1.4 and 1.5 report this measure for Apple, Samsung, and the median non-Apple, non-Samsung firm.[11] As Apple and Samsung roll out cross-device complementarity features, their product lines widen. This pattern motivates the supply-side margin in the model: complementarity can raise the return to carrying products across multiple quality tiers, so product variety may respond through the breadth of the portfolio as well as through the number of models.

[11]The appendix figures provide an alternative min-to-max measure of product variety and illustrate a similar expansionary pattern. Appendix Figure A.1.2 compares the minimum and maximum quality-index values offered by Apple, Samsung, and the industry in each quarter, while Appendix Figures A.1.3, A.1.4, A.1.5, and A.1.6 show the same range-based measure for the main hardware attributes underlying the quality index.

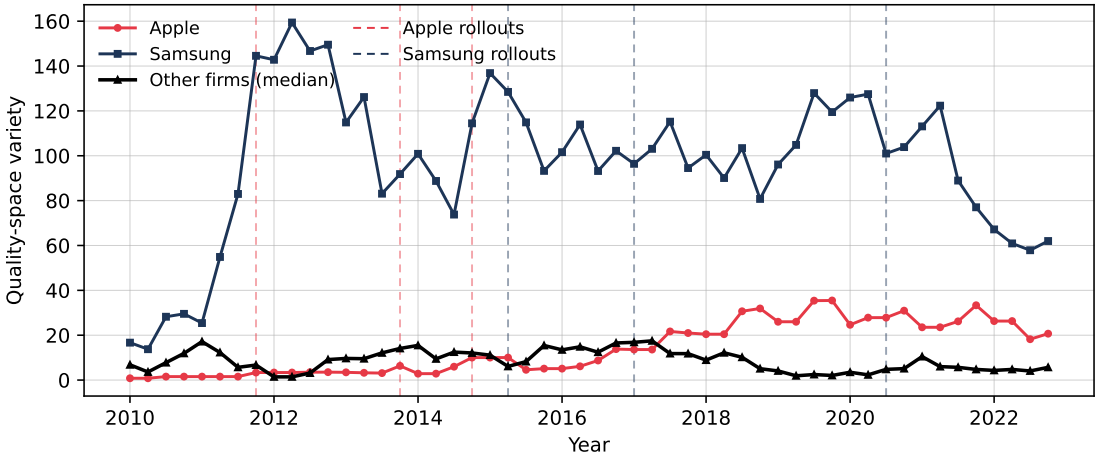


Figure 1.4: Smartphone Quality-Space Variety by Firm Group

Notes: For each quarter and product set, the figure sorts smartphone models by the quality index and computes the Fan–Yang quality-spacing measure as the squared sum of square roots of adjacent quality gaps. The other-firm series reports the median firm-level measure among active non-Apple, non-Samsung firms. All series are plotted on the same scale.

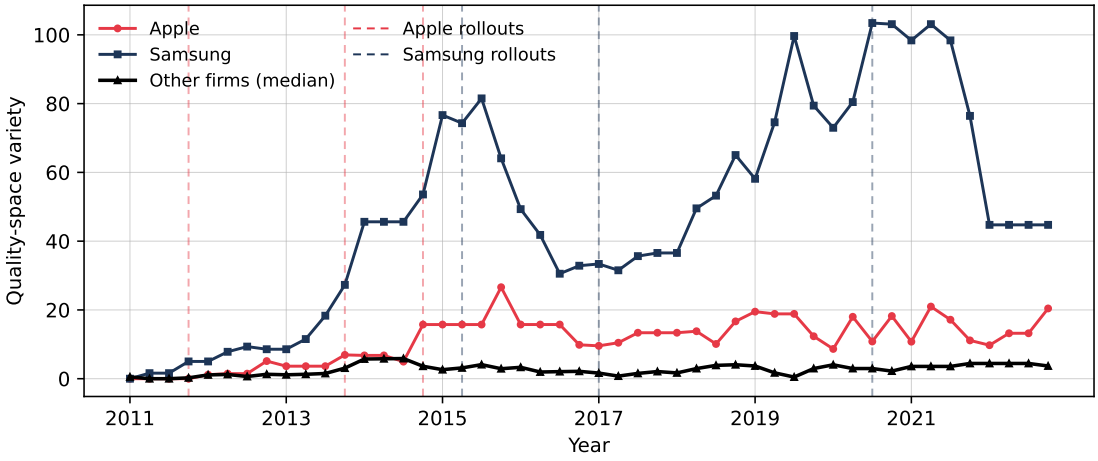


Figure 1.5: Tablet Quality-Space Variety by Firm Group

Notes: For each quarter and product set, the figure sorts tablet models by the quality index and computes the Fan–Yang quality-spacing measure as the squared sum of square roots of adjacent quality gaps. The other-firm series reports the median firm-level measure among active non-Apple, non-Samsung firms. All series are plotted on the same scale.

1.3 Model

1.3.1 Demand

Smartphone and tablet demand follows a random-coefficient discrete choice model. In each period, consumers make purchase decisions separately in the smartphone and tablet markets. Utility from buying a product in one category depends on the consumer's existing device in the other category, so past purchases link demand across categories even though current choices are category-specific. Firms may operate in both markets, set prices, and choose product offerings, taking the current complementarity environment as given.

This paper considers two product categories: smartphones (P) and tablets (T). Let $m \in \{P, T\}$ denote the category of product j, and let $-m$ denote the other category. Individual i in period t owns product d_{it}^{m} in category m and product d_{it}^{-m} in the other category. The indirect utility from purchasing product j is

$$u_{ijt} = \beta_i^{m,0} + \alpha_i^m p_{jt} + \mathbf{x}'_{jt}\boldsymbol{\beta}_i^m + \beta_i^{m,q} q_{jt} + \beta^c c_t(j, d_{it}^{-m}) + \lambda^m \left((t - t_i^m) - \bar{t}^m\right) + FE_{mt} + \xi_{jt} + \epsilon_{ijt}. \tag{1.1}$$

The inside-good intercept varies across consumers according to $\beta_i^{m,0} = \sigma_0^m \nu_i^0$, where $\nu_i^0 \sim \mathcal{N}(0,1)$; this term captures heterogeneity in consumers' baseline propensity to purchase a device in category m in a given period. The price of product j in period t is denoted by p_{jt}. Price sensitivity varies with income according to $\alpha_i^m = \alpha_m + \pi_p^m y_{it}$, where y_{it} is standardized log real income. The vector $\mathbf{x}_{jt}$ contains observed product attributes other than price and quality, including product age, brand dummies, and major-brand trend terms. Preferences for Apple and Samsung products also vary with income: $\beta_{i,b}^m = \bar{\beta}_b^m + \pi_b^m y_{it}$, for $b \in \{\text{Apple}, \text{Samsung}\}$. Product quality is denoted by q_{jt} and is measured using the quality index described in Section 1.2.3. Consumer tastes for quality are given by $\beta_i^{m,q} = \bar{\beta}^{m,q} + \sigma_q^m \nu_i^q$, where $\nu_i^q \sim \mathcal{N}(0,1)$.[12]

Complementarity between product j and consumer i's existing device in the other category, d_{it}^{-m}, enters through the index $c_t(j, d_{it}^{-m}) = \mathbf{x}_t^c(j, d_{it}^{-m})'\boldsymbol{\theta}^c$. The vector $\mathbf{x}_t^c(j, d_{it}^{-m})$ records whether specific complementarity features are available in period t and mutually

[12]Unobserved tastes are assumed to be perfectly correlated across the smartphone and tablet choice problems for a given consumer. In principle, the cross-category correlation in tastes could be estimated directly; in practice, it is weakly identified given the available data. This assumption is conservative for the estimated complementarity parameter. If the true correlation is less than perfect, the model attributes too much cross-category ownership correlation to persistent tastes, thereby attenuating the estimated role of complementarity.

compatible between product j and d_{it}^{-m}. Since the scale of $c_t(\cdot)$ and β^c is not separately identified, the first element of $\boldsymbol{\theta}^c$ is normalized to one. β^c reflects the utility value of the baseline complementarity feature, while the remaining elements of $\boldsymbol{\theta}^c$ measure the relative importance of other features. The replacement term, $\lambda^m \left((t - t_i^m) - \bar{t}^m\right)$, captures state dependence in upgrade incentives, allowing purchase utility to vary as a consumer's current device becomes older relative to the typical replacement horizon. Specifically, t_i^m is the purchase period of the currently owned device and $\bar{t}^m$ is the average replacement cycle for category m; the term therefore depends on the age of consumer i's current device.[13]

The term FE_{mt} denotes category-specific time fixed effects to capture common changes in demand and seasonality. ξ_{jt} denotes product-market demand shocks. Finally, ϵ_{ijt} is an i.i.d. consumer-product-market idiosyncratic shock that follows a type-I extreme value distribution. The mean utility of the outside option is normalized to zero, so that $u_{i0t} = \epsilon_{i0t}$.

1.3.2 Supply

Supply is modeled as a two-stage game within each period. Firms first choose product portfolios, then set prices after demand and marginal-cost shocks are realized.

At the start of period t, firms observe the current complementarity features $\mathcal{C}_t$, existing product portfolios in both markets, current product characteristics, and the ownership distribution generated by past purchases. This state is the simulated distribution over ownership histories generated by the demand system; the empirical ownership moments discipline that distribution but are not themselves the firms' information set. Firms then choose portfolios from their potential product sets. After portfolio choices, current-period demand and marginal-cost shocks are realized, firms set prices, and consumers make purchases that determine future ownership.

Pricing. In the second stage, firms choose retail prices after observing demand shocks $\boldsymbol{\xi}_t$.[14] Firm n solves:

$$\max_{\{p_{jt}, j \in \mathcal{J}_{nt}\}} \sum_{j \in \mathcal{J}_{nt}} (p_{jt} - c_{jt}) MS_t s_{jt}(\mathbf{p}_t, \boldsymbol{\xi}_t) \tag{1.2}$$

[13] By disciplining the evolution of ownership states over time, this term helps the model match brand-ownership patterns, which determine the installed base over which complementarity operates.

[14] The data contain retail sticker prices of devices, not the potentially contracted prices negotiated between device manufacturers and mobile service providers (e.g., discounted iPhones offered through AT&T plans). The pricing model assumes that any such contractual discounts, if present, are uniform across products and do not vary systematically with complementarity features or other product characteristics. This assumption allows observed retail prices to serve as sufficient statistics for firms' pricing decisions in the second stage.

where MS_t is the market size in period t. Equivalently, let $D_{jt}(\cdot) = MS_t s_{jt}(\cdot)$ denote product-level demand. The pricing first-order conditions imply the standard BLP markup equation. Let $\mathbf{s}_t$ and $\mathbf{p}_t$ collect shares and prices within a product category, and let $\Omega^{\text{firm}}_{jk,t} = 1$ if products j and k are sold by the same firm. Define the markup matrix Δ_t by

$$\Delta_{jk,t} = -\Omega^{\text{firm}}_{jk,t} \frac{\partial s_{kt}}{\partial p_{jt}}. \tag{1.3}$$

Then recovered marginal costs satisfy

$$\mathbf{c}_t = \mathbf{p}_t - \Delta_t^{-1} \mathbf{s}_t. \tag{1.4}$$

Marginal cost is specified as:

$$c_{jt} = \mathbf{x}_{jt} \boldsymbol{\gamma}^m + \gamma_q^m q_{jt} + FE_{mt} + \omega_{jt}. \tag{1.5}$$

Here, $\mathbf{x}_{jt}$ includes observed cost shifters and product characteristics, such as brand dummies and product age. q_{jt} denotes the quality index.

Since consumers make current purchase decisions separately for smartphones and tablets based on ownership inherited from past choices, there are no contemporaneous cross-market demand spillovers in the pricing stage. Smartphone and tablet pricing can thus be solved separately in the second stage.

Product Offerings. Product offerings are the first-stage margin through which complementarity affects supply. At the beginning of period t, firms observe the complementarity environment $\mathcal{C}_t$, existing portfolios, product characteristics, and the ownership distribution inherited from past purchases. Following a common assumption in the endogenous product choice literature (e.g., Eizenberg (2014), Wollmann (2018), and Fan and Yang (2020)), firms choose portfolios before current-period demand shocks $\boldsymbol{\xi}_t$ and marginal-cost shocks $\boldsymbol{\omega}_t$ are realized, but know their distributions. Firm n operating in product market $m \in \{P, T\}$ is endowed with a feasible product set $\mathcal{K}^m_{nt}$ and chooses a subset $\mathcal{J}^m_{nt} \subseteq \mathcal{K}^m_{nt}$ to offer. The firm chooses $\mathcal{J}^m_{nt}$ to maximize expected variable profits across product categories net of per-period fixed costs:

$$\max_{\mathcal{J}^m_{nt} \subseteq \mathcal{K}^m_{nt}} \tilde{\pi}^m_{nt}(\mathcal{J}^m_{nt}, \mathcal{J}^m_{-nt}, \mathcal{J}^{-m}_{nt}, \mathcal{J}^{-m}_{-nt}, \mathcal{C}_t) - \sum_{j \in \mathcal{J}^m_{nt}} FC_{jt}. \tag{1.6}$$

where $\tilde{\pi}_{nt}^{m}$ denotes the expected profit from offering products in market m. Specifically, let $\tilde{\mathcal{J}}_{t'} = (\tilde{\mathcal{J}}_{nt'}^{m}, \tilde{\mathcal{J}}_{-nt'}^{m}, \tilde{\mathcal{J}}_{nt'}^{-m}, \tilde{\mathcal{J}}_{-nt'}^{-m})$ denote the portfolios used to evaluate period t' after applying the period-t deviation. Then

$$\tilde{\pi}_{nt}^{m} = \frac{1}{H} \sum_{t'=t}^{t+H-1} \sum_{\ell \in \tilde{\mathcal{J}}_{nt'}^{m} \cup \tilde{\mathcal{J}}_{nt'}^{-m}} \mathbb{E}_{\xi_{\ell t'}, \omega_{\ell t'}} \left[\left(p_{\ell t'}(\tilde{\mathcal{J}}_{t'}, \mathcal{C}_{t'}) - c_{\ell t'} \right) D_{\ell t'}(\mathbf{p}_{t'}, \tilde{\mathcal{J}}_{t'}, \mathcal{C}_{t'}) \right]. \tag{1.7}$$

The profit object is constructed from expected second-stage equilibrium variable profits from offering the products in $\mathcal{J}_{nt}^{m}$ over a horizon of H periods. For each future period, prices and quantities are computed under the induced product portfolios, allowing ownership states to evolve while holding fixed the period-t complementarity feature availability and compatibility structure. Expectations are taken over the distribution of transitory demand and marginal-cost shocks. Age-related product characteristics are updated forward to reflect the decline in demand as products age.

Expected variable profits are aggregated across both product categories. Thus, a deviation in category m affects the product-choice payoff not only through expected profits in that category, but also through the follow-on demand it induces in category $-m$ as ownership states evolve. In forming these expectations, firms take rivals' current product portfolios in both categories as given, and future product offerings are held fixed over the horizon.[15]

The per-period fixed cost of offering product j in category m and period t is parameterized as

$$FC_{jt} = \mathbf{x}_{jt}^{FC} \boldsymbol{\theta}_{m}^{FC} + \nu_{jt}^{m}, \qquad \mathbb{E}[\nu_{jt}^{m} \mid j \in \mathcal{K}_{nt}^{m}] = 0. \tag{1.8}$$

The vector $\mathbf{x}_{jt}^{FC}$ contains product and portfolio characteristics that shift the cost of maintaining product j in the choice set. These include the quality index q_{jt}, product-age bins, brand indicators, and the distance between product j's model quality and the nearest model already offered by the same firm. The quality-distance measure allows fixed costs to vary with the extent to which a product expands the firm's existing line, reflecting additional design, marketing, distribution, or software-support costs associated with products that are farther from the current portfolio.

Observed product portfolios are outcomes of Nash equilibria in product offerings and imply revealed-preference restrictions on fixed costs. Given rivals' product portfolios and the complementarity environment, the observed portfolio must yield weakly higher expected

[15] This restriction is imposed for tractability. The model therefore captures the direct effect of complementarity on product-choice payoffs, while abstracting from dynamic strategic interactions in product offerings across firms and over time.

profit net of fixed costs than all feasible one-product deviations. For notational convenience, define

$$\Pi_{nt}^{m}(A) \equiv \tilde{\pi}_{nt}^{m}(A, \mathcal{J}_{-nt}^{m}, \mathcal{J}_{nt}^{-m}, \mathcal{J}_{-nt}^{-m}, \mathcal{C}_t),$$

where A denotes firm n's portfolio in category m, holding the other portfolios and the complementarity environment fixed.

For an offered product $j \in \mathcal{J}_{nt}^{m}$, removing the product should not increase expected profit net of the fixed cost saved:

$$\Pi_{nt}^{m}(\mathcal{J}_{nt}^{m}) - FC_{jt} \geq \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m} \setminus \{j\}). \tag{1.9}$$

Equivalently, the fixed cost of an offered product is bounded above by the expected incremental variable profit generated by that product:

$$FC_{jt} \leq \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m}) - \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m} \setminus \{j\}) \equiv \overline{FC}_{jt}. \tag{1.10}$$

For a feasible product $j \notin \mathcal{J}_{nt}^{m}$, adding the product should not increase expected profit net of the additional fixed cost:

$$\Pi_{nt}^{m}(\mathcal{J}_{nt}^{m}) \geq \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m} \cup \{j\}) - FC_{jt}. \tag{1.11}$$

Thus, the fixed cost of an absent product is bounded below by the expected incremental variable profit it would generate if added:

$$FC_{jt} \geq \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m} \cup \{j\}) - \Pi_{nt}^{m}(\mathcal{J}_{nt}^{m}) \equiv \underline{FC}_{jt}. \tag{1.12}$$

These inequalities provide upper bounds on the fixed costs of offered products and lower bounds on the fixed costs of absent products. The estimation procedure uses these bounds to construct moment inequalities that discipline the fixed-cost parameters, as described in Section 1.4.

1.4 Estimation

1.4.1 Demand and Marginal Cost

I jointly estimate demand and marginal costs by Generalized Method of Moments using both aggregate product-market shares and the micro-moments described in Section 1.4.1.2. A market is defined as a U.S. quarter, and market size is the U.S. population in that quarter.[16] Because complementarity depends on the stock of devices consumers already own, I simulate ownership states from model-implied purchase probabilities before constructing the ownership moments, as described in Section 1.4.1.1 below. Aggregate demand moments follow the mean-utility inversion in Berry, Levinsohn, and Pakes (1995), which recovers demand residuals by matching model-predicted shares to observed product-market shares. I include cost shifters and characteristics of other products in the same market as instrumental variables. Specifically, the cost shifters include lagged exchange rates and storage and RAM input prices, interacted with product-level average storage and RAM. I also include Gandhi and Houde (2019) differentiation instruments, which shift markups by changing the closeness of a product's substitutes along different dimensions.[17] The maintained timing assumption that firms make product-portfolio choices before shocks realize implies that these instruments are orthogonal to product-quarter demand shocks. The micro-moments complement the aggregate moments by disciplining cross-device ownership patterns and income heterogeneity in brand choice and price sensitivity.

1.4.1.1 Ownership State Simulation

Given that complementarity enters the model through the products consumers already own, demand must track ownership stocks as well as current purchase shares. While aggregate product-market data identify purchase flows, they do not reveal individual smartphone and tablet purchase histories or ownership states. To recover the ownership distribution required

[16] Studies of digital device markets, particularly smartphone markets, often define market size as the population divided by the average number of periods between device purchases. For example, if consumers replace their devices every eight quarters, market size would be the U.S. population divided by eight. Since complementarity depends on population-level ownership stocks rather than quarterly replacement flows, I instead use population as market size and model replacement behavior directly through the demand system.

[17] Exchange rates and component prices enter with one-quarter lags. The selected differentiation instruments are based on local product counts within a half-standard-deviation window of the relevant characteristic in the same quarter. For smartphones, these characteristics include pixel density, screen size, processor speed, and network generation; for tablets, they include pixel density, thickness, processor cores, rear-camera count, and screen size. The selected instruments also include counts of rival Apple and Samsung products.

for demand, I simulate ownership states for a fixed panel of simulated individuals over time. Starting from an initial ownership distribution, model-implied purchase probabilities update next-period ownership probabilities. Since the same persistent tastes may affect both past purchases, and hence current ownership states, and current demand, I hold each consumer's taste draws fixed along the simulated history to preserve this selection.[18]

For a candidate parameter vector, the ownership simulation maps model-implied purchase probabilities into the ownership states that enter future utility. Let $s^m_{idt}(\theta)$ denote the model-implied probability that consumer i purchases product d of type m in period t, and let $s^m_{i0t}(\theta)$ denote the outside-option probability. Let $p^m_{idt}(\theta)$ be the probability that consumer i owns device d of type m at the beginning of period t, with $d = 0$ denoting no owned device. Before the first period for market m, I set $p^m_{i0t^m_0} = 1$ and $p^m_{idt^m_0} = 0$ for $d \neq 0$. After solving the choice probabilities in period t, the ownership probabilities carried into period $t+1$ are updated according to

$$p^m_{i0,t+1}(\theta) = s^m_{i0t}(\theta)p^m_{i0t}(\theta), \tag{1.13}$$

$$p^m_{id,t+1}(\theta) = s^m_{idt}(\theta) + s^m_{i0t}(\theta)p^m_{idt}(\theta), \qquad d \neq 0. \tag{1.14}$$

The transition rule treats replacement as a durable-goods choice, where an inside purchase replaces the consumer's previously owned device, while the outside option leaves ownership unchanged. For products outside the current choice set, $s^m_{idt}(\theta) = 0$, so ownership persists only through the outside option. The resulting ownership states enter current utility through the complementarity term and the replacement term, and generate the model-implied ownership-related moments below. In practice, I use multiple ownership paths for each base simulated consumer to reduce simulation noise, holding taste and income draws fixed across paths while varying ownership draws; detailed simulation settings are reported in Appendix A.4.

Appendix A.4 reports a Monte Carlo validation of the ownership-state simulation using the empirical smartphone and tablet product-market skeleton. In each simulated dataset, I estimate mean coefficients on product characteristics, the quality-taste dispersion, and a scalar complementarity coefficient that turns on during the sample. The design tests whether

[18] I initialize ownership states to match observed overall ownership level in the first quarter of the sample, and subsequent purchase flows and ownership moments discipline the simulated evolution of device stocks. The timing of the sample limits the role of the initial ownership distribution. The sample begins near the emergence of tablets and the early expansion of modern smartphones, so much of the installed base forms within the sample rather than being inherited from an unobserved pre-period distribution.

aggregate shares and simulated ownership histories recover complementarity when current demand depends on devices accumulated through past purchases. Table 1.4 summarizes the results. The estimator recovers the quality-taste dispersion closely, and the estimated complementarity coefficient moves approximately one-for-one with the true complementarity scale. The exercise validates the ownership-state mapping under the maintained transition rule.

Table 1.4: Monte Carlo Validation of Ownership-State Simulation

Scale	Reject	True comp.	Estimated comp.	True σ_q	Est. σ_q
0.0	0.10	0.000	-0.016	0.500	0.503
0.5	1.00	0.500	0.518	0.500	0.503
1.0	1.00	1.000	0.984	0.500	0.502
1.5	1.00	1.500	1.485	0.500	0.502

Notes: The table reports 10 Monte Carlo replications for each complementarity scale, with 500 simulated consumers per replication. Complementarity is a scalar index that is switched on in the middle of the sample period. "Reject" is the fraction of replications rejecting the null that the complementarity coefficient equals zero at the 5 percent level. The estimator jointly recovers the standard demand parameters, the scalar complementarity parameter, and the quality-taste dispersion σ_q.

1.4.1.2 Micro-Moments

I supplement aggregate shares with four sets of micro-moments: (1) Statista cross-brand ownership moments, which inform the distribution of joint ownership across devices and brands; (2) Pew ownership rates, which anchor overall smartphone and tablet ownership stocks; (3) Replacement moments, which identify how these stocks evolve over time; and (4) OpenBrand income moments, which identify income heterogeneity in brand preferences and price sensitivity.

The first set uses Statista brand and cross-brand ownership tables. Let τ index the survey period over which a moment is constructed. For phone brand group b and tablet brand group b', I match

$$\Pr(\text{own tablet brand } b' \mid \text{own phone brand } b,\ \text{own a tablet}; \tau).$$

I also match marginal brand-ownership shares for Apple and Samsung phones and tablets,

$$\Pr(\text{own brand } b \mid \text{own type } m; \tau).$$

The marginal brand-ownership shares provide the baseline distribution of phone and tablet ownership by brand. Same-brand conditional shares above these marginals reveal concentration of ownership within ecosystems; together, these moments anchor realistic ownership shares and help pin down complementarity.

The second set of moments matches overall ownership rates from Pew for smartphones and tablets in the survey years,

$$\Pr(\text{own type } m; \tau).$$

These moments discipline the stock of smartphone and tablet ownership over time. Because the complementarity term enters utility through consumers' existing device ownership, the ownership stock determines the share of consumers exposed to cross-device complementarities. Underpredicting ownership in the other category would understate this exposure margin and could lead the model to rationalize aggregate share co-movement through an artificially large complementarity parameter. Matching ownership levels therefore helps separately discipline the extensive margin of device ownership and the strength of ecosystem complementarities conditional on ownership. Appendix Figure A.1.1 shows the corresponding ownership rates over time.

The third set matches replacement cycles,

$$E(\text{years between acquisitions of type } m; \tau).$$

I match the average interval between consecutive device acquisitions to survey-based replacement intervals for smartphones and tablets. These moments discipline the replacement parameters, which affect the persistence of ownership stocks between purchase occasions and in turn determine the set of consumers exposed to cross-device complementarities.

The fourth set uses OpenBrand purchase records to identify income heterogeneity in current purchases. I construct these moments as differences relative to the low-income group, so that they rely on income gradients rather than purchase-share levels that may differ across datasets due to differences in sample coverage and measurement. I first match

differences in Apple and Samsung purchase shares within price tiers across income groups,

$$\begin{aligned}&\Pr(\text{brand} = b \mid \text{purchase type} = m, p \in \mathcal{P}, y \in \mathcal{Y}; \tau) \\ &\quad - \Pr(\text{brand} = b \mid \text{purchase type} = m, p \in \mathcal{P}, y \in \mathcal{Y}^{Low}; \tau).\end{aligned}$$

By comparing brand shares within price tiers, these moments help isolate income-based brand preferences from income-driven heterogeneity in price sensitivity, in a setting where brands differ systematically in prices. Similarly, to identify the income-price interaction, I also match brand-conditional paid-price differences across income groups,

$$\begin{aligned}&E[p \mid \text{purchase type} = m,\ \text{brand} = b, y \in \mathcal{Y}; \tau] \\ &\quad - E[p \mid \text{purchase type} = m,\ \text{brand} = b, y \in \mathcal{Y}^{Low}; \tau],\end{aligned}$$

where the brand set is $\{\text{Apple}, \text{Samsung}, \text{Other}\}$, with Other pooling all non-Apple and non-Samsung brands. The baseline includes 16 brand-by-price-band share differences and 8 brand-conditional price differences across phones and tablets. Appendix A.3 gives the exact moment definitions. Table 1.7 reports selected ownership, income, and replacement moments, and Appendix Table A.1.2 reports the full set of 49 active micro-moments.

1.4.1.3 Identification

Because the product-market data consist of national quarterly markets, the model relies heavily on changes over time in product choice sets, ownership stocks, and feature availability, together with micro-moments that discipline income heterogeneity.

First, heterogeneity in tastes for quality, σ_q, is disciplined by changes in choice sets across markets along the quality index. Table 1.5 summarizes this variation by reporting product turnover and changes in the quarterly distribution of the quality index. As new products enter and older products exit, consumers face different distributions of price, brand, and quality. The resulting substitution patterns identify the extent to which consumers differ in their willingness to pay for higher-quality or lower-quality devices.

Second, heterogeneity in the value of inside goods, σ_0, is disciplined by changes in the number of available products across markets, together with micro-moments that track ownership shares over time. A larger σ_0 concentrates purchases among consumers with high baseline tastes for inside goods, implying more repeat purchasing and a smaller ownership base for a given level of aggregate purchase flows. A smaller σ_0 spreads purchases more

Table 1.5: Choice-Set Variation Along the Quality Index

Panel A. Choice-set size and turnover

Category	Quarters	Mean active products	SD active products	Mean entries	Mean exits
Smartphones	52	192.1	82.9	32.4	30.9
Tablets	48	80.5	24.9	8.1	7.0

Panel B. Variation in the quality distribution

Category	Quarters	SD mean quality	SD minimum quality	SD maximum quality	Mean within quarter quality SD
Smartphones	52	2.43	1.82	2.74	1.57
Tablets	48	1.33	0.75	2.57	1.86

Notes: Products are model-level products in the final IDC estimation samples. Entries and exits are products whose first or last observed quarter occurs in a given quarter; the first sample quarter is excluded from entry averages and the final sample quarter is excluded from exit averages. Panel B follows the descriptive choice-set variation calculation by first computing the quarterly mean, minimum, maximum, and cross-sectional standard deviation of the category-specific PCA quality index across products. The table then reports the standard deviation over quarters of the quarterly mean, minimum, and maximum quality values, and the average within-quarter quality standard deviation.

broadly across consumers, generating higher ownership rates from the same flows. Figure 1.6 illustrates this mechanism by showing how different values of σ_0 imply different ownership paths over time.

Third, the income-gradient micro-moments discipline the income coefficients on price, Apple preferences, and Samsung preferences, denoted by π_p, π_{Apple}, and π_{Samsung}. Replacement-cycle moments for smartphones and tablets discipline the replacement parameter, λ. Since replacement behavior affects the persistence of ownership stocks, these moments also discipline the installed base that enters the complementarity term.

Finally, the complementarity coefficients are disciplined by co-movements in compatible phone and tablet shares, staggered rollout of complementarity features, and cross-device ownership moments. Feature rollouts vary across firms and features, and each feature raises utility only for mutually compatible phone–tablet pairs. This pair-specific timing helps distinguish complementarity from common demand trends, persistent demand for high-quality devices, and income-based brand sorting. Cross-brand-type ownership moments further constrain the implied joint distribution of smartphone and tablet ownership, providing additional information on whether observed same-ecosystem ownership reflects complementarity rather than only persistent sorting into particular brands.

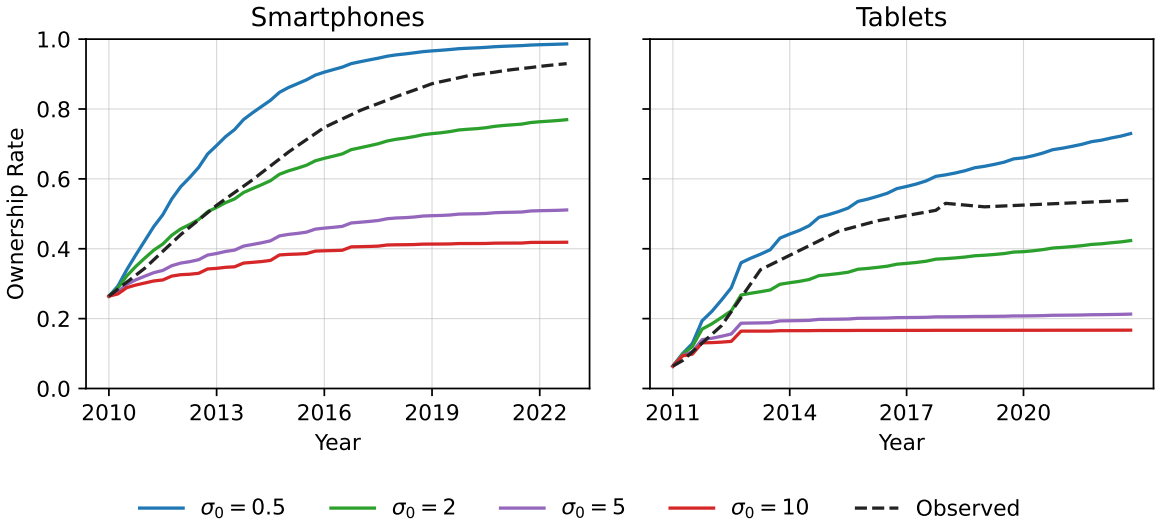


Figure 1.6: Inside-Good Taste Heterogeneity and Ownership Stocks

Notes: The figure maps observed quarterly purchase flows into ownership stocks under alternative values of the inside-good taste dispersion σ_0. For each device category, quarter, and value of σ_0, the exercise solves for a category-level inside-good intercept, analogous to a BLP mean-utility inversion, so that predicted aggregate purchases match the observed aggregate purchase flow in that category. This normalization matches total category purchase flows rather than individual product shares. The first-period ownership stock is normalized to the observed ownership rate. Higher values of σ_0 concentrate purchases among consumers with high baseline tastes for inside goods and therefore generate smaller installed bases from the same aggregate flows.

1.4.2 Fixed Costs

Given the demand and marginal-cost estimates, I recover fixed-cost bounds from the product-choice inequalities in Section 1.3.2. To evaluate product-addition deviations, I construct a set of feasible absent products for each firm, market, and product category. Recently discontinued models provide one source of feasible absent alternatives, but this set is selected on observables: discontinued products are typically older models and may provide limited support for the fixed costs associated with introducing technologically current, differentiated products. I therefore augment the absent-product set with feasible model-level alternatives drawn from the empirical support of product-quality cells to obtain lower-bound inequalities over a broader region of the product space.[19]

Even with supplementing the potential product set with a large set of feasible alternatives, the one-sided product-choice bounds still create a selection problem. Firms observe

[19]Appendix A.5.2 describes the construction of the absent-product set, including the support of feasible quality cells, the minimum quality-distance requirement, and the assignment of non-quality characteristics.

the fixed-cost component ν_{jt}^m when choosing portfolios, while the econometrician does not. Offered products are therefore selected partly on favorable fixed-cost realizations, and absent products are selected partly on unfavorable realizations. In general,

$$\mathbb{E}[\nu_{jt}^m \mid j \in \mathcal{J}_{nt}^m] \neq 0, \qquad \mathbb{E}[\nu_{jt}^m \mid j \notin \mathcal{J}_{nt}^m] \neq 0,$$

even though $\mathbb{E}[\nu_{jt}^m \mid j \in \mathcal{K}_{nt}^m] = 0$ by the fixed-cost specification. Averaging product-drop-deviation upper bounds over offered products, or product-addition-deviation lower bounds over absent products, would therefore condition on selected subsets of the potential set. I therefore construct lower and upper endpoints for every product in the potential set $\mathcal{K}_{nt}^m$. The resulting inequalities average over potential products rather than selected offered or absent products, allowing the fixed cost residual to enter through its unconditional mean-zero restriction. To do this, I make the following assumption.

Assumption 1.1 (bounded fixed-cost support within comparable cells). Within each product category, let $g(j,t,m)$ denote the cell containing product j in period t, defined by year, firm group, model-quality bin, quality-distance bin, and product-age group. Let $\mathcal{A}_g^0$ and $\mathcal{A}_g^1$ denote the absent and offered potential products in cell g, respectively. The direct one-product bounds observed in cell g are

$$\mathcal{D}_g(\theta) = \left\{\underline{FC}_{kt}(\theta) : k \in \mathcal{A}_g^0\right\} \cup \left\{\overline{FC}_{kt}(\theta) : k \in \mathcal{A}_g^1\right\}.$$

Define $V_g^L(\theta) = \inf \mathcal{D}_g(\theta)$ and $V_g^U(\theta) = \sup \mathcal{D}_g(\theta)$. For every feasible product $j \in \mathcal{K}_{nt}^m$,

$$\operatorname{supp}(FC_{jt} \mid j \in \mathcal{K}_{nt}^m,\, g(j,t,m) = g) \subseteq \left[V_g^L(\theta_0), V_g^U(\theta_0)\right].$$

The assumption is similar in spirit to Eizenberg (2014) but restricts the support of fixed costs within narrower cells of comparable products.[20] Since the counterfactual products considered here are device models drawn from the observed product-quality support rather than arbitrary new technologies beyond the current frontier, this restriction is rather mild. Within each cell, products have similar technological sophistication, distance from the firm's existing portfolio, and life-cycle position. The empirical range of one-product profit changes

[20] Eizenberg (2014) imposes a firm-level support restriction, under which fixed costs for all potential products of a firm are bounded by the support of that firm's expected variable-profit changes from single-product deviations. I instead impose the support restriction within characteristic-based cells. This avoids using profit changes from economically distant products and yields more informative bounds for characteristics-based fixed-cost parameters, at the cost of requiring sufficient add- and drop-deviation support within cells.

within the cell therefore provides a conservative support for the missing endpoint.

The completed interval for each potential product is

$$L_{jt}(\theta) = \begin{cases} \underline{FC}_{jt}(\theta), & j \notin \mathcal{J}_{nt}^m, \\ V_{g(j,t,m)}^L(\theta), & j \in \mathcal{J}_{nt}^m, \end{cases} \qquad U_{jt}(\theta) = \begin{cases} \overline{FC}_{jt}(\theta), & j \in \mathcal{J}_{nt}^m, \\ V_{g(j,t,m)}^U(\theta), & j \notin \mathcal{J}_{nt}^m. \end{cases} \tag{1.15}$$

Thus $L_{jt}(\theta_0) \leq FC_{jt} \leq U_{jt}(\theta_0)$ for all potential products. This preserves the revealed-preference content of the add and drop inequalities while replacing the non-randomly missing endpoints with support bounds.

With the completed intervals, I estimate $\boldsymbol{\theta}_m^{FC}$ using the sample analogues of moment inequalities implied by the product-choice restrictions. Let $h_\ell(Z_{jt})$ denote a nonnegative instrument function measurable with respect to information available when the portfolio decision is made. For each product category m and each instrument ℓ, the fixed-cost specification implies

$$\mathbb{E}\left[h_\ell(Z_{jt})\left(\mathbf{x}_{jt}^{FC}\boldsymbol{\theta}_m^{FC} - U_{jt}\right)\right] \leq 0, \tag{1.16}$$

$$\mathbb{E}\left[h_\ell(Z_{jt})\left(L_{jt} - \mathbf{x}_{jt}^{FC}\boldsymbol{\theta}_m^{FC}\right)\right] \leq 0. \tag{1.17}$$

This formulation converts the no-profitable-deviation restrictions into unconditional inequalities, as in the moment-inequality approach to endogenous product offerings. The empirical instruments are indicator functions for bins of model quality, Quality Distance from the firm's existing portfolio, product age, and brand group; for tablets, the specification also uses quality-by-age indicators. The instruments overlap, so a product can contribute to quality-bin, age-bin, and brand-bin moments.

The fixed-cost parameters are set identified. For inference, I follow Andrews, Roth, and Pakes (2023) by inverting coordinate-wise tests based on a profiled max statistic for the linear moment inequalities. I use least-favorable critical values, which yield conservative confidence sets relative to moment-selection refinements but avoid selecting an estimated set of binding product-choice inequalities.

1.4.3 Estimation Results

Table 1.6 reports the joint smartphone–tablet demand estimates. The table first reports mean utility parameters, followed by preference heterogeneity, income interactions, and the common complementarity feature index that links ownership across device categories.

Table 1.6: Demand Estimates

Parameter	Smartphones	Tablets
Mean utility		
Price	-0.0284 (<0.001)	-0.0268 (<0.001)
Quality	2.116 (0.057)	2.012 (0.055)
Model age: 4–7 quarters	-0.255 (0.073)	0.066 (0.064)
Model age: 8+ quarters	-0.296 (0.148)	-0.165 (0.077)
Apple-year trend	0.391 (0.109)	-0.229 (0.075)
Samsung-year trend	0.346 (0.038)	-0.294 (0.023)
Preference heterogeneity and replacement		
S.D. intercept	2.706 (0.004)	2.263 (0.132)
S.D. quality	0.207 (0.004)	0.861 (0.021)
Replacement-age coefficient	0.371 (0.0001)	3.155 (0.104)
Income interactions		
Apple × income	0.249 (0.024)	0.221 (0.031)
Samsung × income	0.108 (0.008)	-0.142 (0.014)
Price × income	0.0016 (<0.001)	0.0056 (<0.001)
Complementarity feature index		
Scale parameter	0.817 (0.030)	
Cloud Sync weight	1	
Share weight	0.289 (0.092)	
Continuity weight	0.345 (0.995)	
Fixed effects		
Year fixed effects	Yes	Yes
Quarter-of-year fixed effects	Yes	Yes
Coarse-brand fixed effects	Yes	Yes

Notes: Cells report coefficient estimates, with standard errors in parentheses. Model age is measured in quarters since model release; the omitted age category is 0–3 quarters. Income is standardized log real income. Fixed effects are absorbed separately by product category. The complementarity term equals the scale parameter times a feature index; Cloud Sync is the normalized feature weight, while Share and Continuity are estimated relative to Cloud Sync.

Table 1.7 reports selected ownership, income, and replacement moments, and Appendix Table A.1.2 reports the full active micro-moment set.

The mean utility estimates show substantial willingness to pay for vertical quality in both categories. Evaluated at the median-income individual with about $52,900 in real annual income, the quality coefficients imply willingness to pay of about $75 for a one-unit increase in the smartphone quality index and $77 for a one-unit increase in the tablet quality index.[21] Because these dollar magnitudes use the income-specific price coefficient, their scale

[21]Dollar conversions divide utility coefficients by the absolute value of the income-specific price coefficient, evaluated at the median-income draw. The median simulated income is about $52,900 in real annual income, corresponding to standardized log income $y_{it} = 0.119$ and price coefficients of -0.0282 for smartphones and -0.0261 for tablets.

Table 1.7: Selected Micro-Moment Fit

Moment	Data	Model
Income sorting		
Phone Apple high-price share, High–Low income	0.134	0.155
Phone Samsung high-price share, High–Low income	-0.113	-0.118
Tablet Apple high-price share, High–Low income	0.065	0.113
Tablet Samsung high-price share, High–Low income	-0.051	-0.103
Phone Apple paid price, High–Low income	57.8	95.3
Phone Other paid price, High–Low income	115.1	77.3
Tablet Apple paid price, High–Low income	41.2	163.2
Tablet Other paid price, High–Low income	61.5	25.7
Cross-device ownership		
Apple tablet share among Apple-phone owners, 2020–2022	0.579	0.647
Samsung tablet share among Samsung-phone owners, 2020–2022	0.392	0.351
Samsung tablet share among Apple-phone owners, 2020–2022	0.122	0.085
Apple tablet share among Samsung-phone owners, 2020–2022	0.224	0.220
Ownership levels		
Apple smartphone owner share, 2019–2022	0.455	0.394
Samsung smartphone owner share, 2019–2022	0.295	0.245
Apple tablet owner share, 2019–2022	0.453	0.393
Samsung tablet owner share, 2019–2022	0.251	0.175
Smartphone ownership rate, 2021	0.850	0.653
Tablet ownership rate, 2021	0.530	0.542
Replacement		
Smartphone replacement cycle	2.71	2.18
Tablet replacement cycle	4.60	4.76

Notes: The table reports selected moments from the active micro-moment set. Data are empirical moments; Model reports fitted values from the estimated model. Share, ownership, and ownership-rate moments are probabilities. Paid-price differences are in real dollars, and replacement moments are in years. Income moments are differences relative to low-income consumers. Appendix Table A.1.2 reports all 49 active micro-moments used in demand estimation.

is linked to the income moments in Table 1.7; the table shows that the model captures the sign of the main high-minus-low income gradients in high-price phone purchases, although the paid-price gradients fit less tightly.

Product age reduces utility more sharply for smartphones than for tablets, consistent with the shorter replacement cycle for smartphones. Relative to products released within the last three quarters, smartphone models aged four to seven quarters are valued about $9 less, and models aged at least eight quarters are valued about $10.5 less. Tablet age discounts are smaller: the four-to-seven-quarter coefficient is close to zero, while the eight-plus-quarter coefficient implies a modest discount of about $6.3.

Brand-specific time trends imply rising valuations for Apple and Samsung smartphones over the sample period, but declining valuations for Apple and Samsung tablets.

Replacement-age coefficients are positive in both categories, so purchase utility rises as

owned devices age relative to the category-specific replacement cycle. Since this term is centered at the average replacement cycle, the larger tablet coefficient reflects a steeper response once tablets are old relative to their longer replacement horizon rather than more frequent tablet replacement overall. The replacement moments in Table 1.7 support this interpretation: the model matches the tablet replacement cycle closely, 4.76 years compared with 4.60 years in the data, while the smartphone replacement cycle is shorter in the model than in the data, 2.18 years compared with 2.71 years.

The estimates also imply substantial preference heterogeneity. The standard deviation of quality tastes is about 10 percent of the mean quality coefficient for smartphones and 43 percent for tablets, implying substantially more dispersion in tablet quality preferences. The standard deviation of the inside-good intercept is also large, suggesting persistent differences in consumers' baseline propensity to purchase devices. This inside-good dispersion helps determine how aggregate purchase flows map into ownership stocks. In Table 1.7, the model closely matches tablet ownership in 2021, 0.542 compared with 0.530 in the data, but underpredicts smartphone ownership, 0.653 compared with 0.850, so the estimated σ_0 should be read together with this remaining stock-level fit gap.

The income interactions are broadly consistent with the purchase-level moments comparing brand choices and paid prices across income groups. Higher-income consumers are less price sensitive in both categories and have stronger preferences for Apple products in both smartphones and tablets. The Samsung-income relationship differs across categories, with higher-income consumers tilting toward Samsung in smartphones but away from Samsung in tablets. For example, the high-minus-low income difference in the Apple high-price phone share is 0.134 in the data and 0.155 in the model; the corresponding Samsung phone moment is -0.113 in the data and -0.118 in the model. The model also matches the signs of the tablet brand-share moments, although it overpredicts the high-price tablet differences and fits the paid-price moments less closely.

Finally, the complementarity estimates imply economically meaningful cross-device demand spillovers. Evaluated at the full complementarity feature index, $0.817 \times (1 + 0.289 + 0.345)$, the estimates imply willingness to pay of about \$47 for a smartphone when the consumer owns a compatible tablet and about \$51 for a tablet when the consumer owns a compatible smartphone. These estimates are supported by a reasonable fit to the cross-device ownership moments in Table 1.7: Apple tablet ownership among Apple-phone owners is 0.647 in the model and 0.579 in the data, and Samsung tablet ownership among Samsung-phone owners is 0.351 in the model and 0.392 in the data. Although not directly identical

in construction, the implied magnitudes are comparable to Lidany (2025), who estimates willingness to pay of about $75 for complementarity features on a compatible smartphone–laptop pair.

Table 1.8 reports local diversion ratios for the top-selling Apple, Samsung, and other-brand products in 2022Q3. These ratios are the most useful price-response summary for the mechanism in this paper because they describe the composition of substitution rather than the level of local price sensitivity. Own- and cross-price semi-elasticities depend on the units of the price change and on baseline product shares; I report them for completeness in Appendix Table A.1.5. By normalizing each cross-price derivative by the corresponding own-price derivative, diversion ratios show where demand goes when a product loses sales.

The diversion ratios make the ecosystem dimension of substitution transparent. A price increase for the iPhone 14 Pro Max diverts more demand to other iPhones than to Samsung or Google models, consistent with stronger substitution within the Apple smartphone line. The selected iPad models also divert strongly to one another. Among tablets, the matrix shows substitution among lower-priced non-Apple products: in particular, the strongest cross-product diversion from the Fire HD 10 within the non-Apple set is to the Galaxy Tab A8. These substitution patterns matter for product choice because they determine how much of a product's lost demand is recaptured by the firm's existing portfolio and how much spills over to rival ecosystems.

Table 1.9 reports marginal-cost estimates used to compute markups, profits, and product-choice payoffs. The recovered first-order conditions imply median margins of about 21.7 percent of price for smartphones and 27.6 percent of price for tablets. The marginal-cost estimates indicate that hardware quality is the main observed per-unit cost shifter: a one-unit increase in the quality index raises marginal cost by about $94 for smartphones and $63 for tablets, holding product age and fixed effects fixed. Relative to products released within the last three quarters, older active models have higher recovered marginal costs, especially among models at least eight quarters old. I interpret these coefficients primarily as controls for systematic age-related variation in recovered per-unit costs among products that remain on the market, rather than as direct evidence that physical production costs rise with age.

Appendix Figure A.1.7 reports the recovered marginal-cost distribution implied by the pricing FOC inversion, separately for smartphones and tablets with Apple and Samsung overlays.

Using the demand and marginal-cost estimates, I recover product-level upper bounds on fixed costs for observed in-market products and lower bounds for feasible but absent prod-

Table 1.8: Demand Diversion Ratios

Panel A. Smartphones									
Price of	13	14 Pro Max	14 Pro	A13	S22	S22 Ultra	Pixel 6a	Moto Power	Moto Stylus
13	–	5.1	4.4	1.2	1.3	1.4	1.5	0.4	0.5
14 Pro Max	7.2	–	5.2	1.0	1.3	1.6	1.3	0.4	0.4
14 Pro	6.9	5.8	–	1.1	1.2	1.5	1.4	0.4	0.4
A13	2.8	1.6	1.5	–	1.7	1.4	1.9	0.8	0.7
S22	3.9	2.8	2.5	2.4	–	2.3	1.9	0.6	0.6
S22 Ultra	4.6	3.7	3.2	2.1	2.4	–	1.8	0.5	0.5
Pixel 6a	3.5	2.2	2.0	2.0	1.4	1.2	–	0.7	0.7
Moto Power	2.6	1.5	1.4	2.1	1.1	0.9	1.9	–	0.7
Moto Stylus	3.1	1.9	1.7	2.0	1.3	1.1	1.9	0.8	–

Panel B. Tablets									
Price of	iPad 2021	Air 2022	mini 2021	Tab A8	S7 FE	S6 Lite	HD 10	HD 8	HD 8 Plus
iPad 2021	–	20.0	6.0	6.8	1.0	0.9	6.2	2.6	2.5
Air 2022	28.6	–	7.6	4.4	1.1	0.7	2.9	0.7	0.7
mini 2021	25.3	22.2	–	4.4	0.9	0.6	3.3	1.0	1.0
Tab A8	14.8	6.8	2.3	–	1.3	1.5	11.0	4.2	3.7
S7 FE	17.8	13.1	3.6	10.7	–	1.4	5.3	1.5	1.5
S6 Lite	16.2	8.4	2.7	12.2	1.4	–	7.4	2.8	2.6
HD 10	13.3	4.3	1.7	10.9	0.6	0.9	–	5.9	4.9
HD 8	8.1	1.4	0.8	6.0	0.3	0.5	8.5	–	7.1
HD 8 Plus	9.3	1.8	0.9	6.3	0.3	0.5	8.5	8.6	–

Notes: Cells report local diversion ratios in percent, evaluated in 2022Q3. Rows index the product whose price increases, and columns index the product receiving diverted demand. An entry is the column product's local share derivative with respect to the row product's price divided by the absolute value of the row product's own-price share derivative. The derivative uses each simulated consumer's individual price coefficient, including the estimated price-by-income interaction, and averages over the income-weighted integration draws. Within each device category, displayed products are the top three Apple products, top three Samsung products, and top three non-Apple/non-Samsung products by market share in the evaluation quarter. Displayed products do not exhaust all diversion destinations. Diagonal entries are omitted.

Table 1.9: Marginal-Cost Estimates

Cost shifter	Smartphones	Tablets
Model age: 4–7 quarters	5.405 (2.443)	11.257 (1.941)
Model age: 8+ quarters	26.933 (4.936)	43.353 (2.006)
Quality	93.650 (0.953)	62.907 (0.609)
Fixed effects		
Year fixed effects	Yes	Yes
Quarter-of-year fixed effects	Yes	Yes
Coarse-brand fixed effects	Yes	Yes

Notes: Cells report marginal-cost coefficients, with standard errors in parentheses. The dependent variable is recovered marginal cost in the same real-price units used in demand estimation. Model age is measured in quarters since model release; the omitted age category is 0–3 quarters. Fixed effects are absorbed separately by product category.

ucts, following the procedure in Section 1.4.2. Figures 1.7 and 1.8 plot these bounds against the quality index, with colors distinguishing Apple, Samsung, and other firms. The left panels report add-deviation lower bounds for discontinued products and feasible hypothetical products; the right panels report drop-deviation upper bounds for observed products. The fixed-cost bounds are higher for Apple and Samsung than for other firms in both categories, but the Apple–Samsung gap is much larger in smartphones than in tablets. This pattern is consistent with the firms' observed smartphone portfolios: Apple offers a narrower set of mostly premium models, while Samsung offers many more models per period, including lower-quality budget devices. The figures show little systematic relationship between quality and the fixed-cost bounds. In smartphones, some same-brand absent products have lower bounds above the upper bounds for observed products with similar quality. This difference reflects product age and other life-cycle characteristics that are not shown in the two-dimensional plots, rather than quality alone.

Using product-level upper and lower fixed cost bounds, I then estimate 95% confidence intervals for fixed cost parameters. Table 1.10 reports projected 95 percent intervals for the fixed-cost parameters for smartphones and tablets.

Table 1.10: Fixed-Cost Parameter Estimates

	Smartphones		Tablets	
Fixed-cost parameter	Lower	Upper	Lower	Upper
Constant	-4.38	20.8	0.27	22.8
Quality	-13.5	-5.39	-3.20	3.65
Quality Distance	-7.10	36.5	-9.92	11.4
Apple	167	282	-9.42	39.2
Samsung	16.9	51.3	-18.0	20.9
Model age: 4–7 quarters	-22.9	1.02	-14.9	-1.29
Model age: 8+ quarters	-40.2	-3.46	-15.8	-3.43

Notes: Entries are in millions of real dollars per model-quarter. "Lower" and "Upper" report the endpoints of the 95% coordinate-wise confidence interval from the profiled least-favorable moment-inequality procedure using diagonal covariance. Quality Distance is the distance between a product's model-level quality index and the nearest product already offered by the same firm in that market. The omitted model-age category is 0–3 quarters.

The coefficients are in millions of dollars per model-quarter and enter the fixed-cost index rather than representing standalone product fixed costs. The constant represents the baseline fixed cost for non-Apple, non-Samsung newly released products, evaluated at zero quality and zero distance from the firm's existing portfolio. Using midpoints of the intervals, this baseline is approximately $8.2 million for smartphones and $11.5 million for

tablets. Conditional on model quality, Quality Distance, and age, the smartphone estimates imply substantially higher fixed costs for Apple and Samsung products than for other phone makers: the Apple interval ranges from $167 million to $282 million, while the Samsung interval ranges from $16.9 million to $51.3 million. This Apple–Samsung ranking is the fixed-cost analogue of the portfolio pattern in Section 1.2: Samsung's greater vertical integration in manufacturing is consistent with lower incremental costs of offering additional smartphone models, whereas Apple's reliance on external manufacturing and narrower premium portfolio is consistent with higher model-level fixed-cost components. Tablet brand intervals are less sharply separated, consistent with the product-level fixed-cost bounds in Figures 1.7 and 1.8. Model age shifts fixed costs downward. At the midpoint, products aged 4–7 quarters lower the fixed-cost index by $10.9 million for smartphones and $8.1 million for tablets, relative to products in their first three quarters; products aged 8 or more quarters lower it by $21.8 million for smartphones and $9.6 million for tablets. The age pattern is consistent with fixed costs being largest around product introduction and declining as design, marketing, distribution, and support costs are amortized over the product life-cycle.

Overall, higher quality does not imply higher fixed costs. For smartphones, the quality coefficient is negative, with a midpoint of about -9.45 million per quality-index unit, and the interval lies below zero. For tablets, the quality coefficient is close to zero, with a midpoint of about 0.23 million and an interval that includes zero.[22] Quality Distance shows a more positive relationship with fixed costs in smartphones: the midpoint is about 14.7 million, although the interval includes zero. This estimate is consistent with differentiated smartphone variants requiring additional product-positioning and distribution resources when they expand the firm's portfolio away from nearby existing models. For tablets, the Quality Distance coefficient is close to zero, consistent with a smaller and less finely segmented product space in which a one-dimensional Quality Distance is less informative about the fixed resources required to maintain an additional model.

1.5 Counterfactual Analysis

Using the estimated model and the 2022Q3 U.S. smartphone and tablet market, I consider three counterfactual complementarity regimes. The first, *no complementarity*, examines

[22]This pattern should not be interpreted as high-quality products being intrinsically cheaper to offer. Rather, conditional on brand, model age, and distance from the firm's existing portfolio, the quality index mainly summarizes hardware content that is strongly reflected in marginal costs, while the fixed-cost index captures per-period design, marketing, distribution, and support costs.

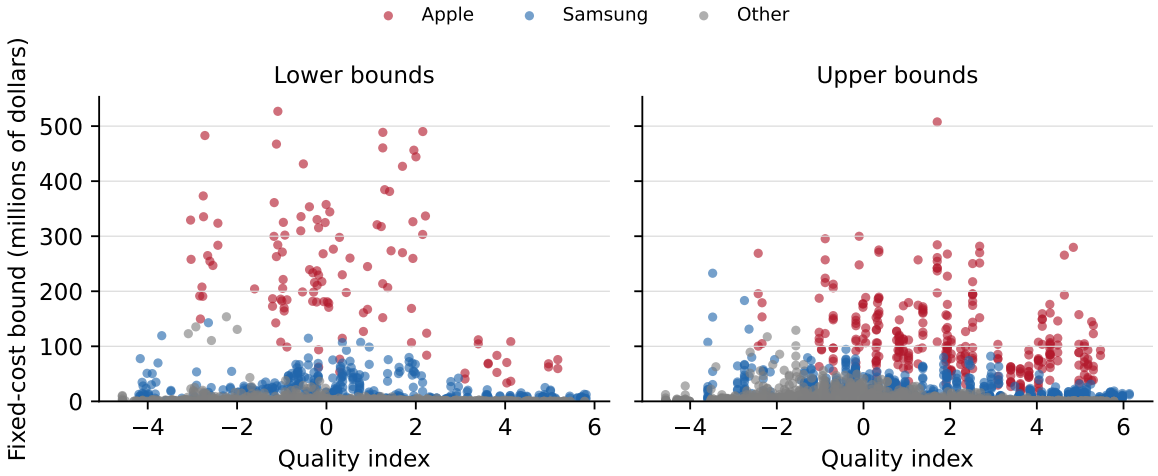


Figure 1.7: Product-Level Fixed-Cost Bounds: Smartphones

Notes: Each point is a smartphone product-market bound directly supplied by a one-product deviation. The left panel plots lower bounds from add deviations for feasible absent products, and the right panel plots upper bounds from drop deviations for observed products. Both panels plot bounds against the model-level quality index. Bounds are in millions of real dollars per model-quarter.

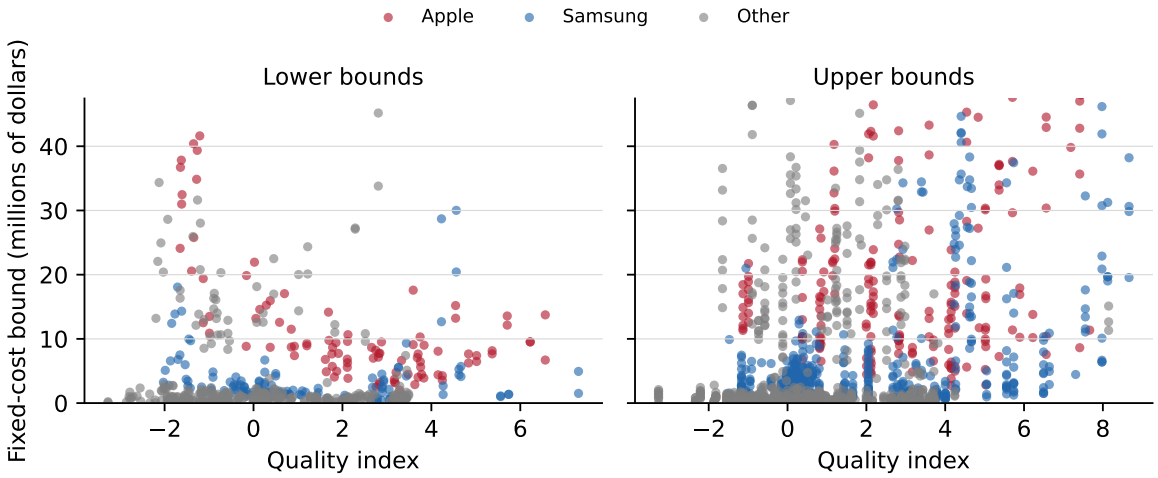


Figure 1.8: Product-Level Fixed-Cost Bounds: Tablets

Notes: Each point is a tablet product-market bound directly supplied by a one-product deviation. The left panel plots lower bounds from add deviations for feasible absent products, and the right panel plots upper bounds from drop deviations for observed products. Both panels plot bounds against the model-level quality index. Bounds are in millions of real dollars per model-quarter.

how ecosystem complementarity affects product variety and welfare. To do this, I remove complementarity by setting Cloud Sync, Share, and Continuity availability to zero for all products. This exercise measures the extent to which observed product variety is sustained by firms' incentives to internalize demand spillovers across their own compatible phones and tablets, and separates the welfare effects of complementarity into three channels: product-variety changes, price changes, and the direct consumer value of cross-device features.

The second counterfactual, *existing-feature interoperability*, examines how interoperability mandates affect product variety and welfare. Specifically, I make active features universally compatible across ecosystems: products that offer a given cross-device feature are treated as compatible with all other products that offer that feature.[23] This policy holds feature availability and estimated preferences fixed, but changes compatibility and therefore which firms can capture the resulting demand spillovers. Under interoperability, a consumer who owns one firm's product can obtain similar complementarity benefits from products in rival ecosystems whenever the relevant cross-device features are active. This weakens the return to products that generate follow-on demand within the firm's own ecosystem, but raises the return to products that become complementary to devices consumers already own from rival ecosystems. Whether the within-ecosystem loss or cross-ecosystem gain dominates depends on installed bases, existing product portfolios, substitution patterns, and fixed costs. The exercise quantifies this net portfolio response and the associated welfare change by recomputing equilibrium product portfolios and prices under existing-feature interoperability.

The third counterfactual, *full-feature interoperability*, is a broader interoperability mandate benchmark. Unlike the existing-feature interoperability counterfactual, which holds observed feature availability fixed and consequently restricts the complementarity effects largely to Apple and Samsung, this exercise allows the interoperability environment to be accompanied by full feature availability: Cloud Sync, Share, and Continuity are turned on for all products and made compatible across all smartphone–tablet pairs. This scenario approximates a more expansive form of interoperability, in which opening cross-device feature interfaces allows other firms to implement the same complementarity features on their own devices. Because the model does not estimate the costs of adopting these features or the pro-

[23] I maintain the observed availability of Cloud Sync, Share, and Continuity, but alter the compatibility matrix so that all devices with a given feature are compatible with each other. In 2022Q3, Cloud Sync, Share, and Continuity are available for the Apple ecosystem and select products in the Samsung ecosystem, while Share is additionally available for newer Android phones and tablets. The counterfactual exercise therefore makes these products mutually compatible within each active feature.

cess by which firms would implement them, I interpret this exercise as a benchmark rather than as a true estimated effect of a broader interoperability mandate; it measures what product variety and welfare would be if interoperability both removed ecosystem exclusivity and made cross-device functionality broadly available.

In each counterfactual regime, I solve the product portfolio game over the same 2022Q3 potential product set.[24] For each firm and product category, the feasible set includes products offered in the baseline market, models discontinued in the previous two quarters, and newly drawn alternatives as constructed in the fixed-cost estimation exercise. Newly drawn products enter with age zero, inherit the firm's brand fixed effect and the period's time fixed effects, and have Quality Distance recomputed relative to the firm's contemporaneous portfolio. Appendix A.5.2 gives the full construction.

Fixed-cost coefficients are set identified, so the counterfactuals require a choice of fixed-cost parameters. Following Wollmann (2018), I evaluate the systematic fixed-cost component at the midpoint of the projected identified set and represent product-level fixed costs as this systematic component, plus a residual. The one-product inequalities used for estimation imply product-specific bounds on these residuals: observed-product drop deviations generate upper bounds, while feasible add deviations generate lower bounds. I use the bounded support restriction in Assumption 1.1 to construct missing sides of these bounds using the empirical distribution of fixed-cost bounds for similar products and draw admissible residual vectors from the resulting product-specific intervals. To isolate policy effects, the same fixed-cost residual vectors are used across the baseline and counterfactual regimes.

For each fixed-cost environment, firms reoptimize product portfolios using the sequential one-product best-response algorithm of Fan and Yang (2020). Starting from the observed 2022Q3 portfolios, firms sequentially consider feasible single-product additions and removals, taking rivals' current portfolios as given, and adopt profitable deviations until no further one-product deviation is available. Within each firm's best-response step, I update the smartphone portfolio first and then update the tablet portfolio conditional on the resulting smartphone portfolio. This ordering reflects the larger smartphone inside share and ownership base, which makes smartphone portfolio choices a natural first step in evaluating cross-device demand spillovers, and is held fixed across all policy regimes. Conditional on the selected portfolio equilibrium, I solve the pricing equilibrium and compute product counts, quality-space variety, prices, consumer surplus, firm profits net of fixed costs, and

[24]For computational tractability, I endogenize product-portfolio choices for the leading firms in each product category: Apple, Samsung, Motorola/Lenovo, LG, ASUS, and Amazon. Portfolios of other firms are held fixed at their 2022Q3 offerings.

total welfare using common shocks across policy environments.

1.5.1 Counterfactual Results

Table 1.11 reports changes in equilibrium outcomes relative to the baseline market. Overall, changes in complementarity and compatibility have economically meaningful effects on product offerings, prices, and welfare, with substantial differences across smartphones and tablets. Removing complementarity lowers total welfare by $1.591 billion, or 17.8 percent relative to baseline, despite a 17.2 percent decline in sales-weighted average prices. Aggregate product count falls by only 1.3 products, or 0.6 percent, and aggregate quality-space variety falls by 38.2 units, or 3.0 percent, but these aggregate changes mask offsetting responses across categories. Both interoperability mandates raise total welfare. Existing-feature interoperability raises total welfare by $0.278 billion, or 3.1 percent, product count by 3.3 products, or 1.6 percent, quality-space variety by 22.7 units, or 1.8 percent, and reduces sales-weighted average prices by 2.8 percent. The broader full-feature interoperability benchmark generates larger effects, raising total welfare by $4.332 billion, or 48.4 percent, with larger changes in product offerings, prices, and surplus.

Table 1.11: Counterfactual Changes in Product Variety, Prices, and Welfare

Policy	Δ Products	Δ Variety	Δ Price	Δ Quality	Δ CS	Δ Profit	Δ Welfare
No complementarity	-1.3	-38.2	-$98.6	-0.61	-$0.808	-$0.783	-$1.591
Smartphones	+0.8	+23.0	-$102.0	-0.63	-$0.726	-$0.756	-$1.482
Tablets	-2.0	-61.3	-$56.2	-0.41	-$0.082	-$0.027	-$0.110
Existing-feature Interoperability	+3.3	+22.7	-$15.9	+0.11	+$0.276	+$0.002	+$0.278
Smartphones	+3.8	+21.2	-$17.1	+0.11	+$0.241	-$0.023	+$0.218
Tablets	-0.5	+1.4	+$4.6	+0.16	+$0.035	+$0.025	+$0.060
Full-feature Interoperability	+11.0	+32.9	-$122.5	-0.98	+$3.000	+$1.332	+$4.332
Smartphones	+9.3	+57.2	-$127.0	-1.01	+$2.818	+$1.224	+$4.042
Tablets	+1.7	-24.3	-$67.7	-0.59	+$0.182	+$0.109	+$0.290

Notes: Entries report changes relative to the baseline 2022Q3 market. Unindented policy rows report aggregate changes; indented rows decompose the same policy change into smartphone and tablet contributions. The rows average across four admissible fixed-cost residual draws. Products are model-level products in the final portfolio. Variety is the Fan–Yang quality-spacing measure; aggregate rows sum the category-specific smartphone and tablet variety changes. Price and quality are sales-weighted averages, using model-predicted equilibrium unit sales as weights, pooled across categories for aggregate rows and computed within category for indented rows. CS, profit, and welfare changes are reported in billions of dollars. Firm profit is net of fixed costs.

Table 1.12 decomposes these welfare changes into three sequential channels: the direct effect of changing the complementarity environment while holding prices and portfolios fixed, the price response holding portfolios fixed, and the remaining effect of endogenous portfolio adjustment. The decomposition shows that the welfare loss from removing complementarity is driven primarily by the direct loss of cross-device feature value, while price adjustment partially offsets this loss. For existing-feature interoperability, the direct compatibility and price-response gains are positive, but they are partially offset by the negative welfare impact from the endogenous portfolio channel. For full-feature interoperability, both the direct complementarity channel and the portfolio channel contribute substantially to the welfare gain.

Table 1.12: Welfare Decomposition by Counterfactual Channel

Policy	Feature/compatibility	Price	Portfolio	Total
No complementarity	-$1.789	+$0.243	-$0.045	-$1.591
Existing-feature interoperability	+$0.878	+$0.067	-$0.666	+$0.278
Full-feature interoperability	+$3.128	-$0.027	+$1.232	+$4.332

Notes: Entries are changes in total welfare relative to the baseline 2022Q3 market, in billions of dollars. The feature/compatibility channel changes the complementarity environment while holding baseline prices and portfolios fixed. The price channel then allows prices to reoptimize holding the baseline portfolio fixed. The portfolio channel is the remaining change from allowing firms to reoptimize product portfolios and associated equilibrium prices. The sequential decomposition is path-dependent by construction.

The remainder of this section discusses the detailed interpretation and firm-level breakdown of each counterfactual exercise.

No complementarity. – Removing complementarity lowers welfare substantially. The welfare decomposition shows that most of the decline comes from the direct loss of cross-device feature value: the feature channel lowers welfare by $1.789 billion, while the price response offsets $0.243 billion of this loss. This price offset is consistent with the model's mechanism: removing complementarity reduces consumers' willingness to pay for compatible devices and weakens firms' pricing incentives, leading to substantially lower sales-weighted average prices. Aggregate product count changes little, but the category decomposition reveals offsetting responses: smartphone offerings rise by 0.8 products, while tablet offerings fall by 2.0 products. Similarly, the aggregate decline in quality-space variety masks an increase in smartphone variety that partially offsets a large reduction in tablet variety. The increase in smartphone offerings and variety reflects an equilibrium replacement effect: removing com-

plementarity weakens the portfolio incentives of firms with established ecosystems, while some rival firms expand into product-space gaps created by these contractions. Table 1.13 illustrates this firm-level reallocation. Apple and Samsung reduce their portfolios, with contractions in both product counts and quality-space variety; the contraction is especially pronounced for Apple tablets, where the portfolio falls from 5 products to 1. Other firms partially offset these exits by expanding their own product offerings. This offsetting entry is nevertheless not enough to fully preserve product count, variety, and welfare. Consumer surplus falls by \$0.808 billion, firm profit falls by \$0.783 billion, and total welfare falls by \$1.591 billion relative to the baseline.

Table 1.13: Product-Portfolio Response Under No Complementarity

Firm	Market	Base products	No-comp. products	Change	Base variety	No-comp. variety	Change
Apple	Phones	10.00	6.50	-3.50	18.2	11.8	-6.4
Apple	Tablets	5.00	1.00	-4.00	13.2	0.0	-13.2
Samsung	Phones	21.00	16.75	-4.25	57.9	50.2	-7.7
Samsung	Tablets	11.00	10.25	-0.75	44.7	48.0	3.2
Others	Phones	13.00	21.50	8.50	35.9	79.4	43.6
Others	Tablets	17.00	19.75	2.75	63.7	68.3	4.6

Notes: Entries report averages across four admissible fixed-cost residual draws for firms whose portfolios are optimized in the counterfactual. Products are model-level products in the 2022Q3 market. Variety is the Fan–Yang quality-spacing measure computed within firm and product market. Others sums the corresponding firm-market entries for all non-Apple, non-Samsung firms. No-comp. denotes the no-complementarity counterfactual.

Existing-feature interoperability. – The existing-feature interoperability counterfactual raises welfare, but the portfolio response is more nuanced than the aggregate increase in product variety suggests. By making active features interoperable across ecosystems, the policy changes not only consumer value and prices, but also the ownership bases over which marginal products generate demand spillovers. This creates two opposing product-portfolio incentives. On the one hand, interoperability weakens the return to products whose profitability depends on generating follow-on demand within the firm's own ecosystem. On the other hand, it expands the set of consumers for whom rival products generate complementarity value, raising the return to differentiated products that appeal to owners of rival devices. In the aggregate, the second force dominates for product counts and quality-space coverage: product count rises by 1.6 percent, and quality-space variety rises by 1.8 percent. Most of this expansion occurs in smartphones, where offerings rise by 3.8 products and variety rises by 21.2 units; tablet offerings fall slightly and tablet variety changes little. Table 1.14 shows

that these aggregate changes mask offsetting firm-level responses. Apple narrows its portfolios in both smartphones and tablets, consistent with the loss of within-ecosystem follow-on demand being more pronounced for a firm with a large installed base and high fixed costs. Samsung expands in both categories, consistent with lower fixed costs and a larger gain from access to the larger share of consumers who own Apple devices. Other firms also expand in smartphones, as select newer products with active Share features become interoperable with Apple and Samsung products, and other products enter to fill gaps in product-space from Apple's contraction.

Table 1.14: Product-Portfolio Response Under Existing-Feature Interoperability

Firm	Market	Base products	Univ. compat. products	Change	Base variety	Univ. compat. variety	Change
Apple	Phones	10.00	7.50	-2.50	18.2	13.5	-4.7
Apple	Tablets	5.00	4.00	-1.00	13.2	9.0	-4.3
Samsung	Phones	21.00	23.75	2.75	57.9	63.4	5.5
Samsung	Tablets	11.00	11.75	0.75	44.7	54.2	9.5
Others	Phones	13.00	16.50	3.50	35.9	43.3	7.4
Others	Tablets	17.00	16.75	-0.25	63.7	61.9	-1.8

Notes: Entries report averages across four admissible fixed-cost residual draws for firms whose portfolios are optimized in the counterfactual. Products are model-level products in the 2022Q3 market. Variety is the Fan–Yang quality-spacing measure computed within firm and product market. Others sums the corresponding firm-market entries for all non-Apple, non-Samsung firms. Existing-feature interoperability keeps feature availability unchanged and makes active features interoperable across ecosystems.

Although existing-feature interoperability increases aggregate product count and quality-space variety, the endogenous portfolio response lowers welfare by \$0.666 billion in the sequential decomposition. This negative portfolio contribution is driven primarily by changes in portfolio composition rather than by induced repricing after portfolio adjustment.[25] The result reflects two features of the existing-feature interoperability experiment. First, the products that exit are valuable: Apple products have high baseline demand and strong estimated consumer value, so their removal can reduce surplus even when rival entry increases the number of products and expands quality-space coverage. Second, feature availability is held fixed; products added by other firms do not offer the full set of cross-device features, so

[25]Holding prices fixed for products that remain in the final portfolio, the change in portfolio composition lowers welfare by \$0.821 billion. Allowing prices to reoptimize after the final portfolio is selected offsets \$0.154 billion of this loss, leaving a total portfolio channel of −\$0.666 billion. The composition loss is concentrated in smartphones: portfolio composition lowers smartphone welfare by \$0.806 billion, compared with \$0.015 billion for tablets.

they do not fully replace the complementarity value associated with the exiting products. As a result, the added or retained products do not fully compensate consumers and firms for the products that leave the market, especially in smartphones. The composition effect also reduces firm profit, indicating that the new portfolio mix is less valuable after accounting for substitution, fixed costs, and demand reallocation.

Prices move in the opposite direction and partially offset the negative portfolio contribution. Interoperability weakens the exclusivity of ecosystem-specific complementarity and reduces the lock-in associated with cross-device features. When consumers can obtain similar cross-device benefits from products in rival ecosystems, ownership of one device no longer makes consumers as captive to the same firm's products in the other category, reducing equilibrium pricing power. In the welfare decomposition, the price response holding portfolios fixed raises welfare by $0.067 billion. Together with the direct compatibility gain of $0.878 billion, this price effect more than offsets the negative portfolio channel. Overall, consumer surplus rises by $0.276 billion, aggregate firm profit is essentially unchanged, and total welfare rises by $0.278 billion relative to the baseline.

Full-feature interoperability. – The full-feature interoperability benchmark generates substantially larger welfare gains than existing-feature interoperability because it expands feature availability in addition to compatibility. Total welfare rises by $4.332 billion relative to the baseline. The welfare decomposition shows that this gain is driven by two forces. First, the direct complementarity channel is large: holding baseline portfolios and prices fixed, turning on Cloud Sync, Share, and Continuity for all products and making each feature interoperable across ecosystems raises welfare by $3.128 billion. Second, unlike in the existing-feature interoperability counterfactual, the endogenous portfolio response is also positive, increasing welfare by $1.232 billion.

The positive portfolio channel reflects a broad expansion in products that become profitable once all firms can offer the same cross-device feature environment. Product count rises by 11.0 products, or 5.5 percent, and quality-space variety rises by 32.9 units, or 2.6 percent. The product-count increase is concentrated in smartphones, where offerings rise by 9.3 products, while tablet offerings rise by 1.7 products. Quality-space responses differ across categories: smartphone variety rises by 57.2 units, while tablet variety falls by 24.3 units. Table 1.15 shows that this aggregate expansion is driven by entry from firms outside the original feature ecosystems. Apple reduces product offerings in both markets, especially tablets, and Samsung's tablet portfolio also narrows. However, these exits are more than offset by substantial expansion from other firms. Unlike in the existing-feature

interoperability counterfactual, the products added by these firms also carry the full set of complementarity features. As a result, rival entry expands product-space coverage while preserving the cross-device feature value that consumers receive from the added products. The positive portfolio channel therefore reflects not only an increase in the number of products, but also a change in the profitability and welfare value of rival products once they participate in the common complementarity environment.

Table 1.15: Product-Portfolio Response Under Full-Feature Interoperability

Firm	Market	Base products	Univ. comp. products	Change	Base variety	Univ. comp. variety	Change
Apple	Phones	10.00	7.33	-2.67	18.2	16.8	-1.4
Apple	Tablets	5.00	1.33	-3.67	13.2	0.6	-12.6
Samsung	Phones	21.00	21.67	0.67	57.9	68.6	10.7
Samsung	Tablets	11.00	9.67	-1.33	44.7	43.3	-1.4
Others	Phones	13.00	24.33	11.33	35.9	108.4	72.5
Others	Tablets	17.00	23.67	6.67	63.7	93.0	29.2

Notes: Entries report firm-market portfolio outcomes for firms whose portfolios are optimized in the counterfactual. Products are model-level products in the 2022Q3 market. Variety is the Fan–Yang quality-spacing measure computed within firm and product market. Others sums the corresponding firm-market entries for all non-Apple, non-Samsung firms. Full-feature interoperability denotes the broader interoperability benchmark defined in Section 1.5.

The comparison between existing-feature interoperability and full-feature interoperability is therefore informative about the range of possible supply-side responses to interoperability. Existing-feature interoperability is a conservative benchmark: it holds feature availability fixed and changes only cross-ecosystem compatibility among products that already offer a given feature. In this case, interoperability raises welfare mainly through direct compatibility gains and lower prices, while the endogenous portfolio response reduces welfare. Full-feature interoperability is an expansive benchmark: it assumes that interoperability is accompanied by full adoption of Cloud Sync, Share, and Continuity across all products. In this case, the portfolio response becomes a major source of welfare gains, and the overall welfare gain is substantially larger. Since the model does not estimate feature-adoption costs or implementation frictions, I do not interpret these exercises as formal bounds on the effect of an interoperability mandate. Rather, they bracket two economically relevant cases: a mandate that only opens existing features, and a broader environment in which opening compatibility also facilitates widespread adoption of cross-device functionality.

1.6 Conclusion

Digital ecosystems create complementarities across product markets: the value of a smartphone can depend on the tablet a consumer owns, and vice versa. This paper studies how these complementarities affect firms' product-portfolio decisions and how this supply-side response changes the welfare consequences of interoperability policy. The central mechanism is that ecosystem complementarity changes the profitability of marginal products. A phone or tablet can generate profit not only through its own sales, but also through the demand it creates for compatible products in the adjacent category. When product portfolios are endogenous, policies that change compatibility therefore affect welfare through direct consumer value, prices, and product variety.

I develop and estimate a structural model of demand, pricing, and product choice in the U.S. smartphone and tablet markets. The demand model allows consumers to value device quality, prices, replacement incentives, brand preferences, and cross-device complementarity with their owned devices. The supply model allows firms to choose product portfolios before setting prices, internalizing demand spillovers across smartphones and tablets. I estimate the model using product-market data from IDC, manually collected information on cross-device feature availability and compatibility, and micro-moments on ownership, replacement, and income-based purchase patterns. The estimates imply economically meaningful willingness to pay for complementarity features and fixed-cost components that differ substantially across firms and product categories.

The counterfactual results show that complementarity supports product variety and affects welfare in ways that would be missed in a price-only analysis. Removing Cloud Sync, Share, and Continuity lowers total welfare by about $1.6 billion in the 2022Q3 market. The direct loss of cross-device feature value is the primary driver, while lower prices partially offset the loss. Product counts change only modestly in the aggregate, but the firm- and category-level responses reveal substantial reallocation: firms with established ecosystems reduce product offerings, while some rival firms expand into the resulting product-space gaps. This pattern illustrates why aggregate product counts alone are insufficient for evaluating the welfare effects of complementarity.

The interoperability counterfactuals highlight a further ambiguity. When observed feature availability is held fixed and active features are made interoperable across ecosystems, welfare rises by about $0.3 billion. The gain comes from direct compatibility benefits and lower prices, while the endogenous portfolio response reduces welfare in the sequential de-

composition. This negative portfolio contribution reflects portfolio composition: the products added or retained under interoperability do not fully compensate consumers and firms for the products that exit, especially in smartphones. In contrast, a broader full-feature interoperability benchmark, in which all products are assumed to offer all cross-device features and these features are interoperable across ecosystems, generates much larger welfare gains. In that exercise, the portfolio response becomes a major source of welfare improvement because rival products also carry complementarity features and become more valuable to consumers outside their original ecosystems.

These results imply that the welfare effects of interoperability depend critically on the supply-side response. Opening compatibility can reduce ecosystem-specific market power and increase consumers' ability to mix products across firms. At the same time, it changes firms' incentives to introduce and maintain products by reallocating cross-category demand spillovers across ecosystems. Whether this reallocation expands or contracts welfare-relevant product variety depends on installed bases, substitution patterns, fixed costs, and whether interoperability also facilitates broader adoption of cross-device features. Policies that focus only on prices or direct compatibility gains may therefore miss an important margin of adjustment.

The counterfactuals do not model the costs or strategic process of adopting new complementarity features; the existing-feature and full-feature interoperability exercises should therefore be interpreted as benchmarks rather than precise forecasts of a particular mandate. Nonetheless, the comparison between them is informative: an interoperability mandate that only opens existing features can have different supply-side consequences from a broader environment in which opening compatibility also makes cross-device functionality widely available. More generally, the results suggest that digital-ecosystem policy should account for both demand-side compatibility gains and the endogenous product-variety incentives created by cross-market complementarities.

References

1. Andrews, Isaiah, Jonathan Roth, and Ariel Pakes. 2023. "Inference for Linear Conditional Moment Inequalities." Review of Economic Studies, 90(6): 2763–2791.

2. Apple Inc. 2024. "Form 10-K for the Fiscal Year Ended September 28, 2024." pp. 5, 7, 46.

3. Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica, 63(4): 841–890.

4. Chatterjee, Chirantan, Ying Fan, and Debi Prasad Mohapatra. 2024. "Spillover Effects in Complementary Markets: A Study of the Indian Cellphone and Wireless Service Markets." Working Paper.

5. Chun, Seungwhan. 2026. "State Dependence and Complementarities in Smartphone Choices." Working Paper.

6. Eizenberg, Alon. 2014. "Upstream Innovation and Product Variety in the U.S. Home PC Market." Review of Economic Studies, 81(3): 1003–1045.

7. Fan, Ying. 2013. "Ownership Consolidation and Product Characteristics: A Study of the U.S. Daily Newspaper Market." American Economic Review, 103(5): 1598–1628.

8. Fan, Ying, and Chenyu Yang. 2020. "Competition, Product Proliferation, and Welfare: A Study of the U.S. Smartphone Market." American Economic Journal: Microeconomics, 12(2): 99–134.

9. Gandal, Neil, Michael Kende, and Rafael Rob. 2000. "The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players." The RAND Journal of Economics, 31(1): 43–61.

10. Gandhi, Amit, and Jean-François Houde. 2019. "Measuring Substitution Patterns in Differentiated-Products Industries." NBER Working Paper 26375.

11. Gentzkow, Matthew. 2007. "Valuing New Goods in a Model with Complementarity: Online Newspapers." American Economic Review, 97(3): 713–744.

12. Lee, Robin S. 2013. "Vertical Integration and Exclusivity in Platform and Two-Sided Markets." American Economic Review, 103(7): 2960–3000.

13. Lee, Sanghak, Jaehwan Kim, and Greg Allenby. 2013. "A Direct Utility Model for Asymmetric Complements." Marketing Science, 32(3): 454–470.

14. Li, Jing. 2019. "Compatibility and Investment in the U.S. Electric Vehicle Market." Working Paper, MIT.

15. Lidany, Yuval. 2025. "Cross-Product Compatibility, Lock-In, and Market Power: The Case of Smartphones and Laptops." Working Paper.

16. Samsung Electronics Co., Ltd. 2024. "2024 Half-year Business Report." pp. 4, 18.

17. Wollmann, Thomas G. 2018. "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." American Economic Review, 108(6): 1364–1406.

CHAPTER 2

Optimal Sin Tax with Endogenous Reformulation

Abstract.[1] This paper studies optimal sugar taxation when firms can reformulate differentiated products. Existing theories of sin taxation typically treat product characteristics as fixed, yet empirically large reductions in sugar intake occur through manufacturers' reformulation responses rather than solely through consumer substitution. We therefore develop a theory of optimal sin taxation with market power and endogenous reformulation, showing that reformulation raises the optimal sugar tax when taxes induce large reductions in sugar content, but lowers it when reformulation is costly or entails large private-utility losses. We then estimate a structural model of demand and supply for the California non-alcoholic beverage market using retail and household-panel data, calibrate reformulation costs using empirical evidence from the U.K. Soft Drinks Industry Levy, and solve optimal-tax counterfactuals. In an efficiency benchmark, incorporating firms' reformulation responses into the counterfactual raises the optimal linear sugar tax from 18 to 48 cents on a 12 g/100ml drink, increases welfare gains from $119.7 million to $784.9 million over a 36-month horizon, and lowers average purchased sugar from 6.27 to 1.74 g/100ml. Although we focus on sugar-sweetened beverages, the framework applies more broadly to corrective taxation in differentiated-product markets where firms can adjust the taxed characteristic.

[1] With Eva Wang and Xinyu Chen.

2.1 Introduction

"Sin taxes," or corrective excise taxes imposed on goods whose consumption results in negative externalities, such as alcohol and tobacco, have been extensively studied in the public finance literature and widely adopted across the world. Recently, taxes on sugar-sweetened beverages have become an area of interest due to rising rates of obesity and cardiovascular diseases associated with excessive sugar consumption. To address this issue, as of 2019, 8 American cities and 39 countries have adopted sugar-sweetened beverage taxes. WHO's 2023 global report documents that such taxes have continued to spread and now take many different forms, including volumetric taxes, ad valorem taxes, and tiered taxes based on sugar content (World Health Organization, 2023).

While much of the literature has focused on the impact of taxes on sin-good prices, empirical evidence suggests that taxes also elicit non-price responses. A prominent example is the Soft Drinks Industry Levy (SDIL) introduced in the U.K. in 2018. The SDIL is a tiered tax on beverages, with rates of 18 pence per litre for products containing between 5 and 8 grams of added sugar per 100ml, and 24 pence per litre for products with more than 8 grams of added sugar per 100ml. In response, over 50% of beverage manufacturers altered the sugar content of their products (HM Treasury, 2018). Moreover, more than 80% of the reductions in sugar intake resulting from the SDIL were due to reformulation activities (Dickson et al., 2025). This points to a powerful channel through which excise taxes reduce consumption of sin goods: when properly designed, sin taxes may induce firms to lower the harmful attributes of their products. More broadly, evidence from multiple U.S. cities finds substantial effects of sugar-sweetened beverage taxes on prices and purchases, while distributional analyses suggest that these taxes can remain progressive on net once revenue use is taken into account (Jones-Smith et al., 2022).

Much of the optimal-tax literature abstracts from this margin. The relevant benchmark for our paper is O'Connell and Smith (2024), who study optimal sugar taxation in differentiated-product markets with market power but treat product characteristics as fixed. Once firms can reformulate, however, a sugar tax can reduce external harm not only by shifting demand away from high-sugar products, but also by inducing manufacturers to lower sugar content directly. Because sugar content rather than beverage volume is the corrective object in this setting, that additional margin can materially change the design of the optimal tax.

Whether reformulation raises or lowers the optimal tax is theoretically ambiguous. Re-

formulation strengthens the corrective case for sugar taxation because part of the reduction in sugar exposure comes through product redesign rather than through a contraction in beverage demand. But reformulation may itself be costly, and lowering sugar can reduce private surplus when consumers value sugar directly. The optimal tax must therefore balance the corrective gains from induced reformulation against the costs of bringing it about.

To study that tradeoff, we develop a framework in which heterogeneous consumers choose among differentiated products, firms with market power choose both prices and sugar content, and the government chooses a tax schedule to maximize social welfare. We then specify and estimate an equilibrium model of supply and demand for the California non-alcoholic beverage market. On the demand side, we estimate a random-coefficients discrete-choice model that allows for heterogeneity in consumer preferences over sugar, price, and other product characteristics. On the supply side, we model a two-stage game in which firms first choose sugar contents, incurring reformulation costs when they alter product composition, and then set prices after observing demand and cost shocks.

The reformulation margin is quantitatively important in the California application. In an efficiency benchmark, a fixed-characteristics counterfactual selects an optimal linear sugar tax of 18 cents on a 12 g/100ml drink. Incorporating manufacturers' reformulation responses raises the optimal tax to 48 cents, increases welfare gains from $119.7 million to $784.9 million over a 36-month horizon, and lowers average purchased sugar from 6.27 to 1.74 g/100ml. Even at the fixed-characteristics tax rate, allowing for reformulation raises welfare gains to $241.5 million, indicating that a policy analysis that abstracts from reformulation misses a substantial share of the gains from sugar taxation.

We also use the framework to study notched sugar taxes as a policy extension. Because notches concentrate incentives around sugar thresholds, they are especially relevant when reformulation is an important margin of response.

This paper contributes to three related literatures. First, it contributes to the public finance literature on optimal sin taxation. Allcott et al. (2019a, 2019b) provide a benchmark framework for optimal soda taxes, while Dubois et al. (2020) and Eichner and Runkel (2025) emphasize the importance of targeting sugar content directly rather than beverage volume. Our paper extends this literature by allowing the tax base itself to respond endogenously through reformulation in a differentiated-product oligopoly.

Second, the paper contributes to the literature on sugar-sweetened beverage taxes. Empirical studies on soda taxes in the U.S. document pass-through, changes in purchases, and reductions in sugar consumption (Cawley and Frisvold, 2017; Cawley et al., 2020; Rojas

and Wang, 2021; Cawley et al., 2021). A growing literature on the U.K. SDIL highlights reformulation as a central policy margin. Dickson et al. (2025) show that reformulation accounts for most of the reduction in sugar intake induced by the levy. Our contribution is to make reformulation an explicit object in the optimal tax problem rather than an auxiliary outcome of policy.

Third, the paper contributes to the literature on imperfect competition and endogenous product choice. O'Connell and Smith (2024) show that market power and redistribution materially affect the design of optimal sugar taxes in differentiated-product markets. We extend that environment by allowing firms to choose sugar content endogenously. In doing so, we also connect to the broader industrial organization literature on endogenous product design and policy-induced changes in product attributes (Fan and Yang, 2020; Fan and Yang, 2022; Barahona et al., 2023; Barwick et al., 2024).

The rest of the paper is organized as follows. Section 2 presents the theoretical framework and the optimal tax problem. Section 3 describes the setup of the empirical model. Section 4 describes the data. Section 5 explains the estimation procedure. Section 6 presents the estimation results. Section 7 calibrates the reformulation-cost parameter using external U.K. SDIL evidence. Section 8 presents the optimal linear sugar-tax counterfactuals. Section 9 concludes.

2.2 Theoretical Framework

We begin with a general optimal-tax condition and then specialize it to a single-rate volumetric tax, a linear tax on the sin characteristic, and a notched tax. Throughout the theory, we use the term sin characteristic to denote the harmful product attribute; in the empirical application, this characteristic is sugar content.

To solve for the optimal tax system, we consider a setting in which individuals choose their consumption of sin goods and substitutes, and the consumption of sin goods is associated with a negative externality imposed on the government in the form of healthcare expenses. We assume that products are differentiated and allow multi-product firms to choose both prices and sin characteristics of their products to maximize profits. The government then chooses tax policies to maximize total welfare, accounting for redistribution and externalities.

2.2.1 Individuals

There is a continuum of individuals, indexed by i, who earn after-tax income Y_i.[2] Individual i earns after-tax labor income z_i and potentially receives a share, $\delta^i \geq 0$, of the after-tax aggregate firm profits $(1-\tau_\Pi)\Pi$, so that $Y_i = z_i + (1-\tau_\Pi)\delta^i\Pi$. Note that $\int_i \delta^i di = 1$.

The market of interest consists of a set of differentiated sin products, a set of alternative goods that serve as substitutes for sin goods, and a composite outside option. $\mathcal{M} = \mathcal{S} \cup \mathcal{N}$ denotes the set of inside goods, i.e., sin products, $\mathcal{S}$ and alternative products, $\mathcal{N}$, and O denotes the composite outside good. We assume that the utility that consumer i derives from consuming good j depends on the quantity of the sin characteristic in the product, denoted by s_j; sin products have $s_j > 0$, whereas other goods have $s_j = 0$. Consumers also face a tax-inclusive price vector $\mathbf{p} = (\mathbf{p}_\mathcal{S}, \mathbf{p}_\mathcal{N}, \mathbf{1})$ for sin goods, alternative goods, and the composite outside option, respectively. The individual's budget constraint can then be summarized by $\sum_{j\in\mathcal{M}} p_j x_j^i + x_o^i = Y^i$.

Individual i chooses consumption of inside goods $\mathbf{x}_\mathcal{M}^i$ given the sin-characteristic vector $\mathbf{s}$ and consumption of the composite outside good x_O^i to maximize individual utility, subject to budget constraints:

$$\max_{(\mathbf{x}_\mathcal{M}^i, x_o^i)} U^i(\mathbf{x}_\mathcal{M}^i, x_O^i, \mathbf{s}, Y^i) \quad s.t. \sum_{j\in\mathcal{M}} p_j x_j^i + x_o^i = Y^i \tag{2.1}$$

From the utility maximization problem, individual i's demand for product j is denoted by $x_j^i = x_j^i(\mathbf{p}, \mathbf{s}, Y^i)$, and individual i obtains indirect utility $V^i(\mathbf{p}, \mathbf{s}, Y^i)$.

We allow utility to depend directly on a product's sin characteristic. This is the key force that gives firms an incentive to choose product characteristics strategically: changing a product's sin characteristic changes consumer demand, and therefore profits. In standard optimal sin tax frameworks, product characteristics are typically taken as fixed, so this channel is absent. In our setting, by contrast, endogenous reformulation is central, and the demand system must therefore depend explicitly on product sin characteristics.

Following the optimal sin tax literature, one could allow consumption to generate both internalities and externalities. In this chapter, we abstract from internalities and focus on external harms, so that the central novel margins are market power and endogenous

[2]We take households' after-tax income as exogenous. Equivalently, the government budget enters the problem through a lump-sum transfer term rather than through interactions with distortionary labor taxation. We therefore abstract from labor-supply responses and from any erosion of labor income tax revenue induced by sin taxes. This allows us to focus on the central tension in our setting, which arises from market power and endogenous reformulation.

reformulation.

2.2.2 Firms

There is a set of firms, indexed by $m = 1, \ldots, M$, where each firm m produces a set of inside goods denoted by $\mathcal{J}_m \subset \mathcal{M}$. For each product j in its portfolio, firm m chooses both tax-inclusive prices, p_j, and sin characteristics, $\{s_j\}_{j \in \mathcal{J}_m}$, to maximize profit, taking into account the market demand for product j, $\mathbf{X}_j(\mathbf{p}, \mathbf{s}, \mathbf{Y}) = \int_i x_j^i di$, the sin tax rate, τ_j, and costs.[3] For each product j offered by firm m, the firm faces a marginal cost, $c_j(s_j)$, where $c_j(s_j) \geq 0$, $c_j'(s_j) > 0$. The firm also faces a fixed cost $FC_j \geq 0$ of offering product j, and a reformulation cost, $RC_j(s_j, \tilde{s}_j) \geq 0$, which depends on both the current sin characteristic s_j and the prior sin characteristic $\tilde{s}_j$. We assume reformulation costs are continuous and increasing in the distance between the new and pre-reformulation sin characteristic. While household income, Y_i, is a function of aggregate firm profits, we assume that firms treat household income as exogenous when making price and sin-characteristic decisions.

We assume that the market is imperfectly competitive, so that firms play a simultaneous-move Nash-Bertrand game over prices and sin characteristics.[4] Firm m's profit maximization problem is given by:

$$\pi_m \equiv \max_{\{p_j\}_{j \in \mathcal{J}_m}, \{s_j\}_{j \in \mathcal{J}_m}} \sum_{j \in \mathcal{J}_m} (p_j - c_j(s_j) - \tau_j) \mathbf{X}_j(\mathbf{p}, \mathbf{s}, \mathbf{Y}) - RC_j(s_j, \tilde{s}_j) - FC_j \tag{2.2}$$

Firm m's optimal choices of $\tilde{p}_j$ and s_j are then functions of marginal costs, reformulation costs, the sin tax rates, market demand for product j, and the price and attribute elasticities of demand. For each product j, the first-order condition for price is given by:

$$\mathbf{X}_j(\mathbf{p}, \mathbf{s}, \mathbf{Y}) + \sum_{j' \in \mathcal{J}_m} (p_{j'} - c_{j'}(s_{j'}) - \tau_{j'}) \frac{\partial \mathbf{X}_{j'}}{\partial p_j}(\mathbf{p}, \mathbf{s}, \mathbf{Y}) = 0 \tag{2.3}$$

Whether a first-order condition exists for the firm's choice of sin characteristic depends on whether RC_j is continuously differentiable. When RC_j is not continuous, such as when a reformulation cost is incurred if and only if $s_j \neq \tilde{s}_j$, then the firm's optimal choice of s_j

[3] At present, the firms' decisions to enter the market and offer products are treated as exogenous in this paper. Given that reduced-sugar or diet versions of the majority of products analyzed here already exist, the benefits of introducing new products are relatively small.

[4] Some studies overlook firm competition details in their theoretical frameworks. However, since product reformulation is a key consideration in this paper, we impose a structure on how firms compete over prices and sin characteristics.

becomes a discrete problem. If RC_j is instead continuous in $(s_j, \tilde{s}_j)$, then for each product j, the first-order condition for the sin characteristic is given by:

$$\Big(-\frac{dc_j}{ds_j}(s_j) - \frac{d\tau_j}{ds_j}\Big)\mathbf{X}_j(\mathbf{p},\mathbf{s},\mathbf{Y}) - \frac{\partial RC_j}{\partial s_j}(s_j,\tilde{s}_j) + \sum_{j'\in\mathcal{J}_m}\Big(p_{j'} - c_{j'}(s_{j'}) - \tau_{j'}\Big)\frac{\partial \mathbf{X}_{j'}}{\partial s_j}(\mathbf{p},\mathbf{s},\mathbf{Y}) = 0 \quad (2.4)$$

To separate the effect of firm-set prices and taxes on tax-inclusive prices faced by consumers, let $\tilde{p}_j = p_j - \tau_j$ denote tax-exclusive prices. Let $\mu_j = p_j - c_j - \tau_j = \tilde{p}_j - c_j$ denote the markup for product j. The aggregate profit in the economy is then characterized by:

$$\Pi \equiv \sum_{j\in\mathcal{M}} \mu_j \mathbf{X}_j - RC_j - FC_j = \sum_{m=1}^{M} \pi_m \quad (2.5)$$

Similar to much of the existing literature on optimal sin taxation (e.g., Allcott et al. (2019b), O'Connell and Smith (2024)), we make a simplifying assumption that the price of the outside option, p_o, is unaffected by any sin tax policies introduced to product set $\mathcal{M}$. This abstracts from substitution toward untaxed goods whose prices may also respond in equilibrium, an issue that may be quantitatively important if such substitution is substantial.

2.2.3 Government

We consider a government that chooses a linear or piecewise-linear sin tax policy to maximize social welfare. Let θ denote the policy instruments, and let $\tau_j(\theta)$ denote the tax liability of product j under policy θ. Given θ, households and firms interact in equilibrium, which determines prices, product characteristics, household income, tax revenue, and external harm.

Let w^i denote the Pareto weight on individual i, and let λ denote the marginal value of government revenue. The government collects revenue from sin taxes and profit taxes. Sin tax revenue from individual i is given by $T_S^i = \sum_{j\in\mathcal{M}} \tau_j x_j^i$, and profit tax revenue is given by $T_\Pi^i = \tau_\Pi \delta^i \Pi$.

Let $\Phi^i(\mathcal{X}_\mathcal{S}^i)$ denote the external harm associated with individual i's consumption of sin goods, where $\mathcal{X}_\mathcal{S}^i \equiv \sum_{j\in\mathcal{M}} s_j x_j^i$ is individual i's total consumption of the sin characteristic. We assume that this externality is increasing in sin-good consumption, i.e. $\frac{\partial \Phi^i}{\partial x_j^i} \geq 0$ for all $j \in \mathcal{S}$, and let $\frac{\partial \Phi^i}{\partial \mathcal{X}_\mathcal{S}^i} \equiv \phi^i$ denote the marginal external harm. We allow ϕ^i to vary across

individuals.

The government chooses θ to maximize the social welfare function

$$\max_{\theta} \int_i w^i V^i(\mathbf{p}, \mathbf{s}, Y^i) + \lambda\Big(T_S^i + T_\Pi^i - \Phi^i(\mathcal{X}_\mathcal{S}^i)\Big) di \tag{2.6}$$

2.2.4 Optimal Sin Tax Policy

In standard sin tax frameworks where sin characteristics are exogenous, externalities from consumption of sin goods are addressed through the price channel only. However, when sin characteristics are endogenous, tax systems with higher tax rates on products with high sin characteristics can lead to greater decreases in externality by incentivizing product reformulations, while maintaining minimal welfare loss in other channels (O'Connell and Smith, 2024; Allcott et al., 2019b; Grummon et al., 2019). To analyze the implications of endogenous reformulation and identify the optimal linear sin tax policy, we consider three different tax systems: a single-rate volumetric tax on sin goods, a linear tax on the sin characteristic, and a single-notch volumetric tax.

2.2.4.1 Optimal Taxation under Smooth Reformulation Costs

We assume reformulation costs are smooth and convex. In the beverage setting, this captures incremental changes in sweetener blends and related ingredients that reduce sugar while preserving sweetness and mouthfeel, with larger departures from the baseline formulation becoming progressively more costly. Under this maintained specification, firms' reformulation choices can be characterized using first-order conditions. We first derive a general optimal-tax condition and then specialize it to the three tax instruments considered below.

We begin with a general, continuous tax system τ_j and associated policy parameters θ. For any equilibrium object a, let $a_\theta \equiv da/d\theta$ denote the total derivative with respect to the policy parameter. The optimal tax policy is characterized broadly by the first-order condition:

$$\int_i \Big(w^i V_\theta^i + \lambda(T_S^i)_\theta + \lambda(T_\Pi^i)_\theta - \lambda(\Phi^i)_\theta\Big)\, di = 0 \tag{2.7}$$

Let $\alpha^i \equiv \frac{dV^i}{dY^i}$, and let $g^i \equiv \frac{w^i \alpha^i}{\lambda}$ represent the social marginal welfare weights. Then the optimal tax policy satisfies:

$$\underbrace{\int_i \left(\frac{dT_S^i}{d\theta} - \frac{\partial T_S^i}{\partial \theta} \right) di}_{\text{fiscal externality}} + \underbrace{\sum_{j \in \mathcal{M}} \left(\mu_j \frac{dX_j}{d\theta} - X_j \frac{dc_j}{d\theta} - \frac{dRC_j}{d\theta} \right)}_{\text{market power distortions}}$$
$$\underbrace{- \int_i \phi^i \sum_{j \in \mathcal{M}} \left(\frac{ds_j}{d\theta} x_j^i + \frac{dx_j^i}{d\theta} s_j \right) di}_{\text{sin externality distortions}}$$
$$\underbrace{+ \int_i \frac{1}{\alpha^i} \left(\sum_{j \in \mathcal{M}} \frac{\partial V^i}{\partial s_j} \frac{ds_j}{d\theta} + (g^i - 1) \frac{dV^i}{d\theta} \right) di}_{\text{distributional concerns}} = 0. \quad (2.8)$$

The optimal policy reflects four forces: fiscal externality, a firm-side term labeled "market power distortions" in the decomposition, sin externality distortions, and distributional concerns. By inducing changes in consumption patterns that affect tax revenue, a change in the sin tax policy can create fiscal externalities. The firm-side term combines the markup channel with induced changes in marginal costs and reformulation costs, while the sin-externality term captures how the policy changes total consumption of the harmful characteristic through both demand responses and reformulation. Finally, the distributional-concerns term captures two channels through which tax policy affects household welfare: the direct utility gain or loss from product reformulation (valued in dollars through $1/\alpha^i$), and the redistribution wedge $(g^i - 1)$ applied to the total utility change, which accounts for differential consumption patterns across individuals with different social welfare weights.

Optimal Single-Rate Volumetric Tax

We begin with a single-rate volumetric tax, $\tau_j = \mathbf{1}(s_j > 0)\tau$, and limit attention to the choice of the common rate τ. For this case, we impose $ds_j/d\tau = 0$, so the tax operates through prices and quantities but not reformulation.[5] The optimal placement of notches is considered later in this section.

Letting $d\mathbb{X}/d\tau \equiv \sum_{j \in \mathcal{S}} \frac{dX_j}{d\tau}$ denote the tax derivative of aggregate demand for sin goods,

[5]This is a modeling assumption rather than an equilibrium implication. It is intended to capture the limited reformulation incentives created by moderate volumetric soda taxes: products far above the zero-sugar threshold face little incentive to adjust sugar content locally, most major products already have zero-sugar versions, and the empirical literature on U.S. volumetric soda taxes finds little reformulation response (Cawley et al., 2020; Cawley and Frisvold, 2017; Rojas and Wang, 2021; Cawley et al., 2021).

specializing equation (8) to $\tau_j = \mathbf{1}(s_j > 0)\tau$ with $ds_j/d\tau = 0$ yields the first-order condition

$$\underbrace{\tau \frac{d\mathbb{X}}{d\tau}}_{\text{fiscal externality}} + \underbrace{\sum_{j\in\mathcal{M}} \mu_j \frac{dX_j}{d\tau}}_{\text{market power distortions}} - \underbrace{\int_i \phi^i \sum_{j\in\mathcal{S}} s_j \frac{dx_j^i}{d\tau} di}_{\text{sin externality distortions}} + \underbrace{\int_i \frac{g^i - 1}{\alpha^i} \frac{dV^i}{d\tau} di}_{\text{distributional concerns}} = 0. \quad (2.9)$$

Rearranging and expanding the distributional-concerns term using $V_\tau^i/\alpha^i = -\sum_j x_j^i(d\tilde{p}_j/d\tau + \mathbf{1}(s_j > 0)) + (1 - \tau_\Pi)\delta^i \, d\Pi/d\tau$, the optimal single-rate volumetric tax is

$$\tau = \underbrace{\frac{\int_i \phi^i \sum_{j\in\mathcal{S}} s_j \, dx_j^i/d\tau \, di}{d\mathbb{X}/d\tau}}_{\text{Pigouvian correction}} - \underbrace{\frac{\sum_{j\in\mathcal{M}} \mu_j \, dX_j/d\tau}{d\mathbb{X}/d\tau}}_{\text{market power correction}} + \underbrace{\frac{1}{d\mathbb{X}/d\tau} \int_i (g^i - 1) \left(\sum_{j\in\mathcal{M}} x_j^i \left(\frac{d\tilde{p}_j}{d\tau} + \mathbf{1}(s_j > 0) \right) - \delta^i (1 - \tau_\Pi) \frac{d\Pi}{d\tau} \right) di}_{\text{distributional correction}}. \quad (2.10)$$

Similar to equation (8), the optimal tax rate balances fiscal externality, sin externality distortions, market power distortions, and distributional concerns. The Pigouvian correction is the standard externality-per-marginal-tax-dollar term, now weighted by individual-specific marginal externality ϕ^i and the sugar content s_j of each sin product whose demand responds to the tax. The market-power correction is typically negative in an imperfectly competitive setting (positive markups, $dX_j/d\tau < 0$ for sin goods), pulling the optimal rate down because taxing an already-distorted quantity amplifies the quantity distortion. The distributional correction vanishes under utilitarian weights ($g^i \equiv 1$) and otherwise captures how price pass-through and profit changes fall on consumers with welfare weights above or below 1.

Optimal Tax on Sin Characteristic

We next consider a tax linear in the sin characteristic, $\tau_j = s_j\tau$. In this case, firms may reformulate, and reformulation costs remain continuous. The tax therefore operates through both demand responses and changes in product characteristics.

Specializing equation (8) to $\tau_j = s_j\tau$ yields the first-order condition

$$\underbrace{\tau \sum_{j\in\mathcal{S}} s_j \frac{dX_j}{d\tau}}_{\text{fiscal externality}} + \underbrace{\sum_{j\in\mathcal{M}} \left(\mu_j \frac{dX_j}{d\tau} - X_j \frac{dc_j}{ds_j}\frac{ds_j}{d\tau} - \frac{\partial RC_j}{\partial s_j}\frac{ds_j}{d\tau}\right)}_{\text{market power distortions}}$$
$$- \underbrace{\int_i \phi^i \sum_{j\in\mathcal{S}} \left(\frac{ds_j}{d\tau}x_j^i + s_j \frac{dx_j^i}{d\tau}\right) di}_{\text{sin externality distortions}}$$
$$+ \underbrace{\int_i \frac{1}{\alpha^i}\left(\sum_{j\in\mathcal{M}} \frac{\partial V^i}{\partial s_j}\frac{ds_j}{d\tau} + (g^i - 1)\frac{dV^i}{d\tau}\right) di}_{\text{distributional concerns}} = 0. \quad (2.11)$$

Unlike equation (9), τ enters the distributional-concerns term implicitly through $dV^i/d\tau$: because $d\tau_j/d\tau = s_j + \tau\, ds_j/d\tau$, the consumer's tax burden changes not just mechanically but also through reformulation at rate τ. Collecting the τ-linear pieces from the distributional term together with the fiscal externality defines the effective sugar base

$$\mathcal{B}_\tau \equiv \sum_{j\in\mathcal{S}} s_j \frac{dX_j}{d\tau} + \int_i (1 - g^i) \sum_{j\in\mathcal{M}} x_j^i \frac{ds_j}{d\tau}\, di,$$

where the first term captures the aggregate sugar-weighted demand response and the second adjusts for the welfare-weighted reformulation-induced change in tax burden. Solving for τ yields the implicit function:

$$\tau = \underbrace{\frac{\int_i \phi^i \sum_{j\in\mathcal{S}} \left(\frac{dx_j^i}{d\tau}s_j + \frac{ds_j}{d\tau}x_j^i\right) di}{\mathcal{B}_\tau}}_{\text{Pigouvian correction}} - \underbrace{\frac{\sum_{j\in\mathcal{M}} \left(\mu_j \frac{dX_j}{d\tau} - X_j \frac{dc_j}{ds_j}\frac{ds_j}{d\tau} - \frac{\partial RC_j}{\partial s_j}\frac{ds_j}{d\tau}\right)}{\mathcal{B}_\tau}}_{\text{market power correction}}$$
$$+ \underbrace{\frac{1}{\mathcal{B}_\tau}\int_i (g^i - 1)\left(\sum_{j\in\mathcal{M}} x_j^i \left(\frac{d\tilde{p}_j}{d\tau} + s_j\right) - \delta^i(1 - \tau_\Pi)\frac{d\Pi}{d\tau}\right) di}_{\text{distributional correction}}$$
$$- \underbrace{\frac{1}{\mathcal{B}_\tau}\int_i \frac{g^i}{\alpha^i} \sum_{j\in\mathcal{M}} \frac{\partial V^i}{\partial s_j}\frac{ds_j}{d\tau}\, di}_{\text{reformulation-utility correction}}. \quad (2.12)$$

The optimal tax on the sin characteristic balances the same four forces as equation (8): fiscal externality (now scaled by $\mathcal{B}_\tau$ rather than a simple demand elasticity), a firm-side

correction, sin externality distortions, and distributional concerns. Relative to O'Connell and Smith (2024), which treat product characteristics as fixed, the Pigouvian correction now picks up two channels through which the tax reduces sin consumption: standard demand substitution ($dx_j^i/d\tau$) and product reformulation ($ds_j/d\tau$). As in O'Connell and Smith, market power implies a pre-existing quantity distortion, so the term labeled market-power correction typically pushes the optimal tax down relative to the competitive benchmark; in the present setting that term should be interpreted broadly, because reformulation affects both marginal costs ($dc_j/ds_j \cdot ds_j/d\tau$) and reformulation costs in addition to markups. The distributional and reformulation-utility corrections separate: the former handles how price pass-through and profit changes redistribute across consumers, while the latter accounts for the direct welfare impact of altered product characteristics on consumers who have preferences over s_j. When reformulation is endogenized, $ds_j/d\tau$ enters nearly every component.

The central implication is therefore conditional. When the externality from sugar is large and reformulation is an effective way to reduce sugar exposure, allowing for endogenous reformulation strengthens the corrective case for taxation, so the optimal sugar tax can exceed the fixed-characteristics benchmark in O'Connell and Smith (2024). By contrast, when reformulation is costly, or when induced reductions in s_j generate substantial private-surplus losses, this additional margin is less valuable, and the optimal tax is correspondingly lower; in the limit, as reformulation costs become prohibitive, the problem approaches the fixed-characteristics benchmark. Put differently, the sign and magnitude of the difference relative to O'Connell and Smith depend on how efficiently the tax induces manufacturers to reduce sugar content rather than simply incur reformulation costs. This logic also motivates the notched-tax exercise below: when the externality is large and much of the welfare gain comes from discrete reductions in sugar content, a notch may outperform a linear sugar tax by concentrating incentives at salient reformulation thresholds.

Optimal Notched Tax

We next consider a notched tax of the form $\tau_j = \bar{\tau}\mathbf{1}(s_j > \underline{s})$. Relative to a linear tax on the sin characteristic, a notch concentrates incentives on products near the threshold. In a differentiated-products oligopoly with endogenous reformulation, however, the products that cross the threshold are not known a priori: they are determined jointly by demand, multiproduct pricing, and reformulation costs. The notched-tax problem is therefore an equilibrium policy problem rather than a local sufficient-statistic problem.

Numerical Mechanism Exercise

We study this problem in a normalized four-product duopoly and compare a linear sugar

tax, $\tau_j = \tau s_j$, with a single notched tax, $\tau_j = r\mathbf{1}(s_j > \underline{s})$. The exercise is intended to clarify when threshold-based incentives to reformulate can dominate the more uniform incentives created by a linear tax. Table 2.1 reports the maximal welfare gain attained by the best linear tax and the best notch in each supply environment. Appendix B.1 provides the primitives and computational details.

Table 2.1: Toy Mechanism Exercise: Welfare Comparison of Linear and Notched Taxes

Supply environment	Linear sugar tax	Notched tax	Equilibrium response
Fixed characteristics	0.084	0.074	Prices adjust, but sugar content is fixed. The linear tax is slightly better because it targets sugar intensity directly.
Smooth convex reformulation cost	0.146	0.172	The notch sets $\underline{s} = 3.0$ and induces two products to cross the threshold. The linear tax induces all four products to reformulate smoothly.

Notes: Entries report the maximal welfare gain under each tax instrument, relative to the no-tax allocation, in the normalized toy economy. Welfare is measured as consumer surplus plus producer surplus plus tax revenue minus external harm and any reformulation costs. The notched tax is a per-volume tax paid only by products with $s_j > \underline{s}$; the zero-threshold case is nested in this instrument.

Table 2.1 highlights the central mechanism. With fixed characteristics, the linear tax performs slightly better because it targets sugar intensity directly. With smooth reformulation costs, by contrast, the notch can outperform by concentrating incentives on products near the threshold and inducing discrete reductions in sugar content while avoiding some of the output distortion from taxing inframarginal sugar on products that remain below the cutoff.

2.3 Empirical Model

2.3.1 A Structural Approach

Much of the existing optimal-tax literature expresses the optimal tax formula as a function of elasticities, sufficient statistics, or other objects that can be calibrated using empirical estimates. In the present setting, however, that approach becomes inadequate. As O'Connell and Smith (2024) argue, the optimal tax formula is influenced by firm-side distortions,

which require knowledge of markups and marginal costs that are hard to observe directly. Endogenous reformulation adds a second difficulty: the policy problem depends on how firms adjust product characteristics in response to the tax schedule, and empirically estimating those responses in a setting with heterogeneity and strategic interaction is challenging. We therefore estimate an equilibrium model of supply and demand for the California non-alcoholic beverage market. For tax systems with analytic solutions, the model delivers the elasticities and firm-side margins needed to implement the formulas. For tax systems that do not admit an analytic characterization, it provides the computational environment for comparing candidate tax schedules in equilibrium.

2.3.2 Demand

We estimate demand for non-alcoholic beverages following the standard random-coefficient discrete-choice approach proposed by Berry, Levinsohn, and Pakes (1995). Household i derives indirect utility from product j in market m and month t according to

$$u_{ijmt} = \alpha_i p_{jmt} + \beta_i^{sugar} sugar_{jt} + \gamma^{sugar^2} \frac{sugar_{jt}^2}{10} + \beta_i^{juice} x_j^{juice} + \beta_i^0 + \zeta_{b(j)mt} + \xi_{jmt} + \epsilon_{ijmt} \quad (2.13)$$

where p_{jmt} is the price of product j per 20-oz equivalent, $sugar_{jt}$ is grams of sugar per 100ml, x_j^{juice} equals 1 if product j is a juice drink, and $\zeta_{b(j)mt}$ stacks brand-county, brand-month, retailer-county, retailer-month, county-month, and brand-retailer fixed effects. We retain the quadratic sugar term, scaled by ten, to capture diminishing marginal utility in sugar. The term ξ_{jmt} is an unobserved product-market-month demand shock and ϵ_{ijmt} is an i.i.d. type-I extreme-value taste shock. We include an outside option indexed by $j = 0$ and assume that, in each month, the household selects the alternative that maximizes utility. The monthly choice frequency mitigates concerns about consumer stockpiling behavior.[6]

Consumer heterogeneity enters through two random coefficients and four demographic interactions:

$$\begin{aligned}
\alpha_i &= \bar{\alpha} + \alpha_y y_i \\
\beta_i^{sugar} &= \bar{\beta}^{sugar} + \beta_y^{sugar} y_i \\
\beta_i^{juice} &= \bar{\beta}^{juice} + \beta_y^{juice} y_i + \sigma^{juice} \nu_i^{juice} \\
\beta_i^0 &= \beta_y^0 y_i + \sigma^0 \nu_i^0,
\end{aligned}$$

[6]Hendel and Nevo (2006) document stockpiling at the weekly frequency in the U.S. soft-drinks market.

where y_i is standardized log household income and $\nu_i^{juice}, \nu_i^0 \sim N(0,1)$ are independent unobserved-heterogeneity draws. Since sugar taxes operate through prices that may depend on sugar content, the model needs to capture how preferences for sugar consumption covary with price sensitivity. We generate this correlation by allowing both sugar preferences and price sensitivity to depend on household income.[7]

2.3.3 Supply

We assume that beverage manufacturers play a two-stage static game. We abstract away from firm-retailer relationships and assume that manufacturers set consumer prices. In doing so, we assume that manufacturer-retailer markups do not change in counterfactual simulations.

In stage 1, firms simultaneously choose the sugar content of existing products. In stage 2, they play a simultaneous-move Nash-Bertrand game in which they choose retail prices after observing demand and marginal cost shocks. This timing assumption is standard in the endogenous-product-choice literature and follows Eizenberg (2014) and Fan and Yang (2020).

- **Stage 2**: In month t, after observing demand and marginal cost shocks $(\xi_{jmt}, \omega_{jmt})$, firm n chooses p_{jmt} to maximize

$$\max_{p_{jmt},(j)\in\mathcal{J}_{nmt}} \sum_{(j)\in\mathcal{J}_{nmt}} (p_{jmt} - mc_{jmt}) D_{jmt}(\mathbf{p}_{mt}, \mathbf{sugar}_t) \tag{2.14}$$

  where the marginal cost is

$$mc_{jmt} = \eta^0 d_{jt}^{diet} + \eta^{sugar} sugar_{jt} + \eta^{juice} x_j^{juice} + \eta^{sugar\times juice}(sugar_{jt} \cdot x_j^{juice}) + \psi_{b(j)mt} + \omega_{jmt}. \tag{2.15}$$

  Here $d_{jt}^{diet} = \mathbf{1}(sugar_{jt} = 0)$ is a zero-sugar indicator, and $\psi_{b(j)mt}$ stacks the same six sets of two-way fixed effects used on the demand side. The parameters estimated in this stage, beyond the fixed effects, are $(\eta^0, \eta^{sugar}, \eta^{juice}, \eta^{sugar\times juice})$.

- **Stage 1**: At the beginning of quarter T, before observing demand and marginal cost shocks $(\xi_{jmt}, \omega_{jmt})$, firm n chooses the sugar content of products in its portfolio to

[7] Because random coefficients enter only the juice indicator and the inside-good intercept, within-product substitution across retailers inherits the IIA structure of the type-I extreme-value shock. This restriction is less important for the sugar-tax counterfactuals in this paper, which operate primarily through changes in sugar content and cross-product substitution.

maximize expected variable profits net of reformulation costs:

$$\max_{sugar_{jT}, j \in \mathcal{J}_{nmT}} \sum_{j \in \mathcal{J}_{nmT}} \left(\pi_{njT}(\mathbf{sugar}_T) - RC_j(sugar_{jT}, sugar_{j,T-1}) \right) \tag{2.16}$$

where

$$\pi_{njT}(\mathbf{sugar}_T) = \sum_{t \in \mathcal{T}_T} \mathbb{E}_{\xi_{jmt}, \omega_{jmt}} \left\{ \Big(p_{jmt}(sugar_{jt}, sugar_{-jt}) - mc_{jmt}(sugar_{jt}) \Big) D_{jmt}(\mathbf{p}_{mt}, \mathbf{sugar}_t) \right\} \tag{2.17}$$

π_{njT} is the expected stage-2 T-period profit from product j given sugar-content vector $\mathbf{sugar}_T$. The empirical implementation uses a continuous quadratic reformulation cost,

$$RC_j(sugar_{jT}, sugar_{j,T-1}) = \kappa Q_j \left(\frac{sugar_{jT} - sugar_{j,T-1}}{10} \right)^2, \tag{2.18}$$

where sugar is measured in grams per 100ml and Q_j is baseline product quantity, measured in 20-oz equivalents, over the payoff horizon. Since the California data do not contain policy-induced reformulation, κ is calibrated using external evidence from the U.K. SDIL rather than estimated from observed California reformulation choices.

One limitation of the empirical model is that it does not observe the use of non-caloric sweeteners, such as aspartame and sucralose, in either demand or supply. In practice, manufacturers often offset sugar reductions with such ingredients in order to preserve sweetness and mouthfeel. Because nutrition labels do not report the quantity of these sweeteners, the model treats reformulation as changes in sugar content alone. This omission may cause the recovered marginal costs of lower-sugar products to reflect the cost of substitute ingredients, and it may also overstate the demand response to sugar reductions in counterfactuals if firms can preserve taste with unobserved sweetener substitution.

2.4 Data

We study the California non-alcoholic beverage market in calendar year 2018, focusing on carbonated beverages and juice drinks. The empirical analysis combines Nielsen Retail Scanner data, which provide store-week prices, quantities, brand, size, and category information, with Nielsen Consumer Panel data, which provide household purchase behavior and

demographics. We aggregate the retail data to the monthly level and use sales-weighted average prices within month. Because neither Nielsen source reports nutrition information, we match products to the USDA Branded Food Data and Open Food Facts using UPCs and hand-collect sugar content for the remaining unmatched UPCs among the top 20 best-selling brands. County demographics come from the American Community Survey, and the input-cost series used to construct price instruments come from Federal Reserve Economic Data.

The unit of observation in the estimation is a county-month market. We define a product as a brand-category-sugar content-retailer quadruple, so UPCs with the same brand, product category, sugar content, and retailer are aggregated to a common product. Sizes are converted to 20-oz equivalents, and product prices are computed as volume-weighted averages across sub-UPCs within market-month. A product is treated as available in a market-month if it belongs to a brand with at least 0.5% market share among all brands in the Nielsen RMS data in 2018. Under this definition, the estimation sample accounts for roughly 72% of all sales in the Nielsen Retail Scanner data.

The resulting sample spans 648 county-month markets and 3,693 product-market observations produced by 16 firms across 61 brands and sold through 33 retail chains. We define market size as the median monthly non-alcoholic beverage sales in a county, measured in 20-oz equivalents, divided by 0.16, the average monthly share of household trips on which a beverage is purchased in the panel data. The household panel serves two roles in the analysis. First, it provides the micro moments that discipline heterogeneity in demand. Second, it describes purchase patterns in the category. Table 2.2 shows that households typically undertake 2 trips on which beverages are purchased per month and buy 2 unique products.

Table 2.2: Average Monthly Consumption, Trip, and Retailer Statistics

	Mean	25^{th} pct.	Median	75^{th} pct.	SD
Number of Products	3.21	1	2	4	3.07
Number of Trips	2.34	1	2	3	1.81
Number of Retailers	1.58	1	1	2	0.90
Number of 20-oz Units Purchased	21.68	5.07	12.67	28.36	27.76

* This table provides summary statistics for the monthly number of unique non-alcoholic beverages purchased, monthly number of trips on which beverages were purchased, number of retailers visited in a month, and the monthly total 20-oz equivalent units purchased per household in the Nielsen Consumer Panel data in 2018.

Table 2.3 summarizes the market-level structure of the estimation sample. The average

county-month market contains 117 products and 43 brands, and the in-sample products account for about 72% of total scanner volume.

Table 2.3: Product, Brand, Market-Share, and Volume Statistics by Market

	Mean	Median	SD
Number of Products	117.22	118	16.88
Number of Brands	43.11	43	1.22
Share of In-sample Products	0.7244	0.7254	0.0371
Number of 20-oz Units Sold Per Month	730195.5	277228.3	1531110
Number of County-Month Markets	648		

* This table presents market-level statistics of the unique number of products, brands, share of in-sample products, and units sold per month. Share of in-sample products reflects the share of total beverage volume in the Nielsen data accounted for by in-sample products in 2018, and does not represent the market shares used in estimation.

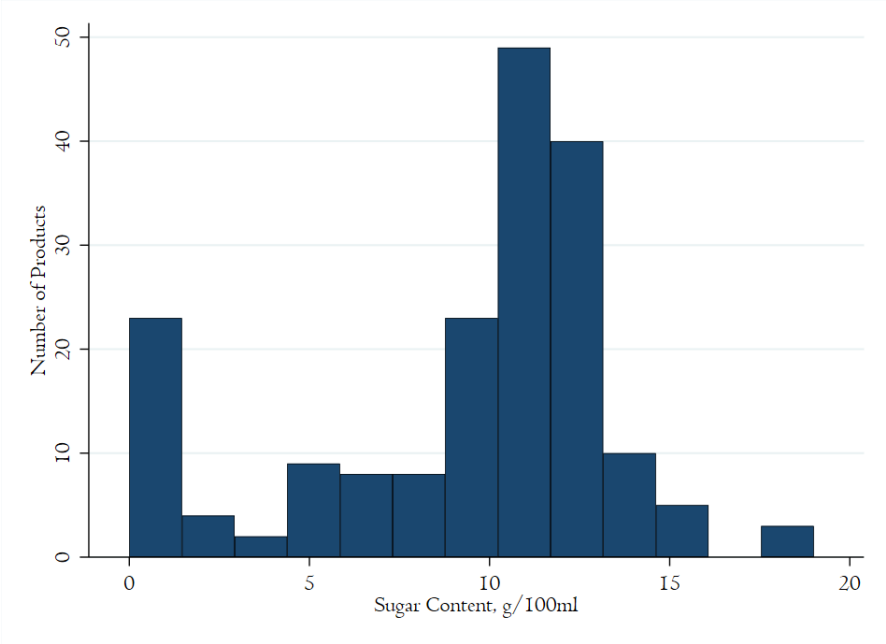
Table 2.4 summarizes the in-sample products by category and sugar status. The sample is concentrated in sugar-added products, especially juice drinks, while zero-sugar products appear only among carbonated beverages. Figure 2.1 displays the distribution of sugar content for all products in the sample. The distribution exhibits a rough bimodal pattern, with diet products primarily located at a sugar level of 0, while sugar-added products cluster in the range of 11 to 13 grams per 100ml.

Table 2.4: Product, Brand, Market-Share, and Price Statistics by Product Category

Product Category	Sugar Content	# Products	# Brands	Share (%)	Price
Carbonated Beverages	With sugar	51	16	31.72	0.8128
	Zero sugar	13	12	17.97	0.7622
Juice Drinks	With sugar	133	17	24.14	0.8173
All		197	45	73.83	0.8020

* This table contains the unique number of products and brands in each product category. Shares reflect the relative proportion of total beverage volume accounted for by in-sample products in 2018, and do not represent actual market shares used in estimation. Prices are weighted by sales across products.

Figure 2.1: Sugar Content of In-Sample Products, 2018



2.5 Estimation

Following Fan and Yang (2022) and Petrin (2002), we use three sets of moments to identify the demand and supply parameters.

The first set of moments are macro moments. Following Berry, Levinsohn, and Pakes (1995), we match observed market shares to model-predicted market shares to identify the mean taste coefficients and the absorbed fixed effects.

To address potential price endogeneity, we instrument price with variables constructed from input costs: 1) the monthly average U.S. price of sugar interacted with product sugar content, 2) the monthly average U.S. price of corn-based sweetener interacted with product sugar content, 3) the monthly global price of aluminum interacted with a product's sales-weighted average size, and 4) the producer price index of plastic bottles interacted with a product's sales-weighted average size. We also interact these variables with a product-category dummy to allow for cost heterogeneity across product types.

The second set of moments uses household-panel micro data to discipline heterogeneity. These moments target expenditure per inside purchase by income bin, the number of inside purchases per month overall and by income bin, the number of juice-drink purchases per month overall and by income bin, and total sugar consumed through purchased beverages

per month by income bin.

The third set of moments uses the correlation of product choices with income to identify the income-taste coefficients. Estimation is carried out by two-step GMM with clustered weighting matrices, where demand and marginal costs are jointly estimated.

2.6 Estimation Results

Table 2.5 reports demand estimates. Demand for non-alcoholic beverages meaningfully price sensitive, with grouped own-price elasticities ranging from -2.41 to -1.82. The estimates also imply a natural pattern of sugar valuation: consumers value sweetness, but the marginal value of sugar declines with sugar content. At the mean coefficients, corresponding to a household with average income, the sugar component of utility peaks at about 7.8 grams per 100ml.

Table 2.5: Demand Estimates

Mean Coefficients (β)	price	-3.40	Income Interactions (π)	price	0.064
		(1.01)			(0.034)
	sugar	0.239		sugar	-0.007
		(0.056)			(0.006)
	sugar^2/10	-0.154		juice	-0.024
		(0.044)			(0.055)
	juice	-1.207		intercept	0.118
		(0.239)			(0.067)
Random Coefficients (σ)	juice	1.649			
		(0.060)			
	intercept	1.558			
		(0.077)			

Standard errors clustered at the brand level are in parentheses. Income interactions are reported in standardized log-income units (mean zero, unit standard deviation across simulated agents). The linear sugar coefficient varies across households only through the income interaction. An earlier specification included a free random coefficient on the linear sugar term, disciplined by the sugar analog of the juice micro moment used in the estimation. This moment was matched well at zero through income-based heterogeneous sugar taste, so we restrict sugar taste heterogeneity to enter only through the income interaction. Brand–county, brand–month, retailer–county, retailer–month, county–month, and brand–retailer fixed effects are absorbed.

The estimates also imply economically sensible heterogeneity. Variation across house-

holds through unobserved taste shocks is concentrated in tastes for juice drinks and in baseline participation in the packaged-beverage category, rather than in large differences in sugar tastes.[8] Higher-income households are less price sensitive, more likely to choose inside goods, less likely to purchase juice, and have weaker preferences for sugar. For a product with 5 grams of sugar per 100ml, the implied marginal willingness to pay for one additional gram of sugar is 2.64 cents per 20-ounce equivalent for households with an annual income of roughly $21,000, compared with 2.31 cents for households with an annual income of roughly $149,000.

Table 2.6 reinforces this interpretation. Cross-price elasticities are positive but modest, suggesting that products are substitutes without implying unrealistically narrow diversion patterns. A 1 percent increase in the prices of all sugar-sweetened beverages lowers sugar-sweetened-beverage volume by 1.94 percent.[9] The model therefore implies that sugar taxes operate through both reductions in aggregate sugar-sweetened beverage purchases and substitution across beverages, the two margins central to the policy analysis. These magnitudes are also close to the disaggregated cola elasticities reported by O'Connell and Smith (2024), who find own-price elasticities around -2 for selected Coke products and Coke–Pepsi cross-price elasticities of about 0.02–0.04 for comparable package sizes.

[8] An earlier specification included a free random coefficient on the linear sugar term, disciplined by the sugar analog of the juice micro moment used in the estimation. This moment was matched well at zero through income-based heterogeneous sugar taste, so we restrict sugar taste heterogeneity to enter only through the income interaction.

[9] For comparison, O'Connell and Smith (2024) estimate a corresponding elasticity of 1.41 using U.K. data.

Table 2.6: Pooled Grouped Own- and Cross-Price Elasticities

	Sugar (g/100ml)	(1)	(2)	(3)	(4)	(5)	(6)
(1) Coca Cola	11	-1.94	+0.04	+0.02	+0.04	+0.02	+0.05
(2) Coca Cola Diet	0	+0.08	-2.04	+0.02	+0.03	+0.02	+0.05
(3) Coca Cola Zero	0	+0.08	+0.04	-2.41	+0.03	+0.02	+0.05
(4) Pepsi	12	+0.08	+0.03	+0.02	-1.82	+0.02	+0.05
(5) Sprite	11	+0.09	+0.04	+0.02	+0.04	-2.15	+0.05
(6) Gatorade Perform 02	6	+0.05	+0.02	+0.01	+0.03	+0.01	-2.18

Entries report the percent change in total row-group quantity from a 1 percent common price increase for all products in the column group. Product-level elasticities are aggregated across county-month markets using baseline 20-oz-equivalent quantities.

Groups aggregate brand-product-retailer alternatives across retailers, FIPS, and months. The two other-soda groups and the juice-drink group are selected by pooled baseline quantity.

Pooled row quantities (millions of 20-oz equivalents): (1) 232.4, (2) 103.1, (3) 39.4, (4) 113.9, (5) 61.5, (6) 175.8.

Markets with row-group availability: (1) 648, (2) 648, (3) 646, (4) 648, (5) 648, (6) 648.

Table 2.7 reports the marginal cost estimates from the supply side, recovered jointly with demand under the multi-product Nash-Bertrand assumption. Conditional on the absorbed fixed effects, the zero-sugar intercept is significantly negative, implying that diet products carry lower recovered marginal costs than otherwise comparable sugar-added products. Among non-juice products, the negative sugar coefficient implies that lower observed sugar is associated with higher recovered marginal cost, while the positive interaction $\eta^{sugar \times juice}$ implies that this relationship is much weaker for juice drinks. These coefficients are best read as reduced-form cost slopes within the pricing system rather than as literal engineering cost parameters. Because the model does not observe non-caloric sweeteners or other taste-preserving ingredients, any cost differences associated with greater reliance on those inputs in lower-sugar products are loaded onto the sugar dimension. For the reformulation counterfactuals, this is the relevant object: if reducing sugar requires firms to substitute toward more expensive non-caloric sweeteners or other taste-preserving inputs, the estimated sugar-cost slope captures the resulting change in marginal cost and therefore preserves firms' incentives to reformulate.

Table 2.7: Marginal Cost Estimates

	Estimate
zero-sugar indicator (η^0)	−0.587
	(0.176)
sugar (η^{sugar})	−0.046
	(0.013)
juice (η^{juice})	−0.417
	(0.186)
sugar × juice ($\eta^{sugar \times juice}$)	0.036
	(0.017)

Standard errors clustered at the brand level are in parentheses. The same six sets of two-way fixed effects used in the demand specification are absorbed in the supply side as well.

Table 2.8 compares the fourteen relative micro moments used in estimation to their model-implied counterparts. The moments split into four families: expenditure per inside purchase by income bin, the number of inside purchases per month overall and by income bin, the number of juice-drink purchases per month overall and by income bin, and total sugar consumed through purchased beverages per month by income bin. The model fits most moments closely.

Table 2.8: Micro Moments: Observed vs. Model-Implied

Moment	Observed	Model-Implied	Difference
Expenditure per inside purchase, high income	0.700	0.741	−0.041
Expenditure per inside purchase, middle income	0.676	0.724	−0.048
Expenditure per inside purchase, low income	0.646	0.712	−0.066
Number of inside purchases per month	9.207	8.712	+0.495
Number of inside purchases, high income	9.648	9.267	+0.382
Number of inside purchases, middle income	9.361	8.590	+0.771
Number of inside purchases, low income	8.256	8.429	−0.173
Number of juice purchases per month	5.103	4.998	+0.104
Number of juice purchases, high income	5.386	5.201	+0.185
Number of juice purchases, middle income	4.961	4.849	+0.111
Number of juice purchases, low income	4.663	5.024	−0.360
Sugar amount per month, high income	67.370	70.629	−3.260
Sugar amount per month, middle income	74.344	67.277	+7.066
Sugar amount per month, low income	65.672	67.892	−2.221

Moments are computed from the Nielsen Consumer Panel for California households in 2018 ($n = 4{,}886$ household-month observations). Income bins use raw annual income thresholds of \$45,000 and \$120,000. Expenditure per inside purchase is measured in dollars per 20-oz equivalent; number of inside purchases and number of juice purchases are counts of purchase occasions per household per month; sugar amount per month is total grams of sugar per household per month consumed through purchased beverages. Each moment is normalized by its observed value when entering the GMM objective.

Figures 2.2 and 2.3 provide a final supply-side diagnostic. They plot the distributions of implied markups, $p_{jmt} - mc_{jmt}$, and marginal costs mc_{jmt} pooled across all 429,165 product–market observations. The median markup is \$0.31 per 20-oz equivalent and the median Lerner index $(p_{jmt} - mc_{jmt})/p_{jmt}$ is 0.34, indicating substantial but not extreme market power in the California non-alcoholic beverage sector. The lower tail of the markup distribution lies close to the single-product Bertrand floor of $1/|\beta_p| \approx \$0.29$ implied by $\beta_p = -3.40$.

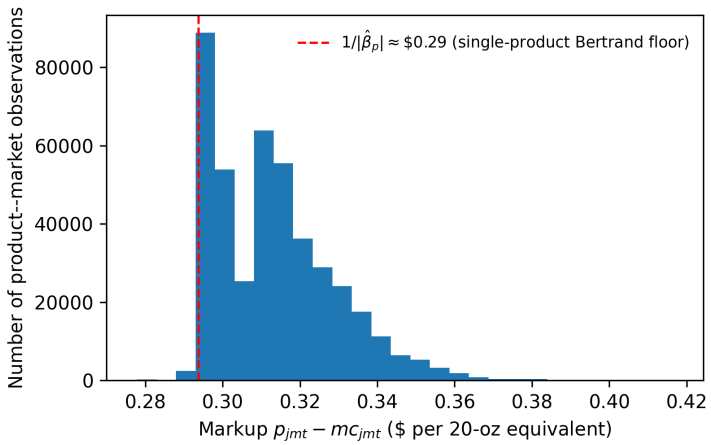


Figure 2.2: Distribution of implied markups

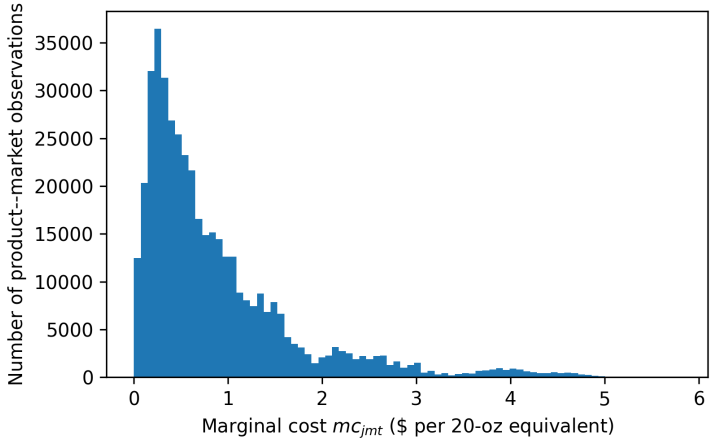


Figure 2.3: Distribution of implied marginal costs

2.7 Reformulation Cost Calibration

A central empirical challenge is that reformulation costs are not observed. In principle, observed reformulation choices could set-identify these costs: a product that reformulates must have reformulation costs below the gain from doing so, while a product that does not reformulate must have costs high enough to make reformulation unprofitable. The California sample, however, contains too few reformulations to make these inequalities informative; this is expected in a setting without a policy shock that creates sharp incentives for firms

to move products across sugar thresholds. We therefore use external policy evidence from the U.K. Soft Drinks Industry Levy (SDIL) to discipline the reformulation-cost parameter that governs firms' endogenous sugar responses in the counterfactuals. The SDIL generated substantial reformulation at well-defined sugar thresholds, taxing liable drinks at 18 pence per litre when sugar content is at least 5 and below 8 grams per 100ml, and at 24 pence per litre when sugar content is at least 8 grams per 100ml.

The calibration sets the common reformulation-cost parameter that would make the estimated California model reproduce the product-level reformulation rate implied by the SDIL evidence, under the same tax policy. For each eligible formulation j with baseline sugar $s_j > 5$, we compute the present value of the unilateral firm-portfolio profit gain from changing only product j's sugar content to a candidate value s'_j:

$$PV\Delta\Pi_j(s'_j) = \sum_{h=1}^{H} \delta_h \left[\Pi_{f(j),h}(s'_j, s_{-j}; \tau^{SDIL}) - \Pi_{f(j),h}(s_j, s_{-j}; \tau^{SDIL})\right], \tag{2.19}$$

where $\Pi_{f(j),h}$ is the variable profit of the firm that owns product j in horizon h, including all products in the firm's portfolio in the markets where j is sold. Both terms in equation (2.19) are post-tax profits: the comparison is between keeping the observed sugar content under the SDIL and reformulating under the same SDIL schedule. After each unilateral sugar change, prices are re-solved in the stage-2 Nash-Bertrand game and marginal costs are updated using the estimated supply equation. This profit object lets the firm internalize cannibalization across its own products without requiring a full simultaneous reformulation equilibrium.

The cost of moving from s_j to s'_j is

$$RC_j(s'_j, s_j) = \kappa Q_j \left(\frac{s'_j - s_j}{10}\right)^2, \tag{2.20}$$

where Q_j is baseline 20-oz-equivalent quantity over the same payoff horizon used in equation (2.19). The quantity term means that κ is measured in dollars per baseline 20-oz equivalent for a 10 g/100ml sugar-content change over the payoff horizon. Scaling reformulation costs by quantity allows the calibration to accommodate the U.K. pattern in which some high-volume flagship products remained taxed while many other products moved below the thresholds.[10] The parameter should be interpreted broadly: it captures recipe, testing,

[10]A pure product-level sunk cost, independent of sales, would have to be calibrated high enough to rationalize non-reformulation by high-volume products; mechanically, that value would make reformulation prohibitively expensive for smaller products.

packaging, and supply-chain costs, as well as any quality or taste loss from sugar reduction that is not separately modeled through observed product characteristics.

The candidate set follows the SDIL thresholds. Products with $5 < s_j < 8$ may move to 4.9 g/100ml, just below the lower threshold. Products with $s_j \geq 8$ may move either to 7.9 g/100ml, avoiding the high tier while remaining in the low tier, or to 4.9 g/100ml, avoiding the levy entirely. Given a value of κ, product j chooses the candidate that maximizes net value:

$$\hat{s}_j(\kappa) = \arg\max_{s' \in S_j} \left\{ PV\Delta\Pi_j(s') - \kappa Q_j \left(\frac{s' - s_j}{10} \right)^2 \right\}, \tag{2.21}$$

where staying at s_j is always available and has value zero. The model-predicted reformulation rate is

$$R(\kappa) = \Pr\left(\hat{s}_j(\kappa) \neq s_j \mid s_j > 5\right). \tag{2.22}$$

We choose κ so that $R(\kappa)$ matches external evidence on the share of initially levy-eligible products that reformulated. Because the external evidence does not pin down a single transition moment perfectly aligned with the California product universe, we report a target range. The main range, 35–40 percent, is anchored by recent product-level evidence on reformulation under the SDIL and by evidence that the levy caused large movements of products out of taxable sugar bands. Appendix B.2 gives the tax conversion, representative-month scaling, and choice-set details.

Table 2.9: Reformulation Cost Calibration

Target	κ	Count share	Sales share	Below 5	5–8	Mean $\Delta sugar$
30%	0.463	30.2%	26.1%	12.6%	17.6%	-0.72
35%	0.449	35.2%	27.8%	12.6%	22.6%	-0.87
40%	0.444	40.3%	30.5%	12.6%	27.7%	-1.03
50%	0.391	50.3%	31.4%	12.6%	37.7%	-1.43

The eligible set is products with baseline sugar above 5 g/100ml. Count share is the share of eligible reformulation units whose chosen sugar differs from baseline. Sales share weights reformulation units by baseline quantity in the calibration sample. Below 5 and 5–8 report the post-calibration shares of eligible units in each sugar band. Mean $\Delta sugar$ is the count-weighted average change in grams per 100ml across all eligible units. The benchmark uses firm-portfolio profit gains, re-solved post-tax prices, a 36-month horizon, and the last observed month across all California FIPS as a representative monthly profit environment.

Table 2.9 reports the calibrated values. The benchmark uses the last observed month across all California FIPS as a representative monthly profit environment, a 36-month payoff horizon, a 5 percent annual discount rate, and re-solves post-tax prices for each unilateral deviation. Under the central 35–40 percent target range, the calibrated κ lies between 0.444 and 0.449 dollars per baseline 20-oz equivalent for a 10 g/100ml sugar-content change. At these values, the model predicts that 27.8–30.5 percent of eligible sales volume reformulates and that the count-weighted average sugar reduction among eligible products is 0.87–1.03 g/100ml. The average present-value gross profit gain among predicted reformulators is \$1.22–\$1.31 million, and the corresponding average implied reformulation cost is \$330,000–\$360,000. In this sense, the calibration implies a reformulation margin that is quantitatively important but far from costless.

This calibration should be interpreted as a benchmark for counterfactual policy analysis rather than as a structural estimate identified from California reformulation choices. It uses U.K. evidence to discipline a California model, and it evaluates unilateral rather than simultaneous reformulation incentives.[11] Because simultaneous competitor responses would lower each firm's marginal benefit from reformulating, the calibrated κ should be viewed as an upper bound on the equilibrium value that would rationalize the same target reformulation rate in a full stage-1 game. Appendix B.2 provides the remaining implementation details.

2.8 Optimal Linear Sugar Taxes

We next use the estimated model to quantify how endogenous reformulation changes the design and welfare consequences of an optimal linear tax on sugar concentration under an efficiency benchmark.[12] The section compares three cases. First, we solve the fixed-characteristics benchmark, in which firms keep observed sugar contents and adjust only prices after the tax is introduced; this is the natural benchmark relative to O'Connell and Smith (2024). Second, we evaluate that same tax when firms are allowed to reformulate, which isolates the welfare value of the reformulation margin itself. Third, we solve for the

[11] Although applying U.K. SDIL reformulation evidence to a California model is imperfect, it provides a useful external calibration target because the estimated California demand system implies elasticities close to those estimated in U.K. data.

[12] Welfare is the money-metric change in consumer surplus, producer surplus, tax revenue, and externality savings, net of reformulation costs, evaluated over a 36-month horizon with a 5 percent annual discount rate. We set $g^i = 1$ for all consumers, $\lambda = 1$, and the profit-tax wedge to zero. External harm is calibrated following the O'Connell-Smith implementation of Wang et al. (2012), so each physical gram of sugar generates 0.7 cents of external damage. Product j pays τs_j per 20-oz equivalent, where s_j is grams of sugar per 100ml, so a 12 g/100ml drink pays 12τ.

optimal tax when firms may adjust sugar content before repricing. This sequence shows both how much is missed by shutting down reformulation and how much the optimal tax changes once reformulation is incorporated into the policy counterfactual.

Benchmark

Table 2.10 compares the fixed-characteristics benchmark with our model's optimal-tax recommendation once firms may reformulate before repricing. The first row reports the fixed-characteristics benchmark in the spirit of O'Connell and Smith (2024); the second reports the reformulation-allowed optimum implied by our model. Relative to the fixed-characteristics benchmark, the reformulation-allowed optimum features a higher tax rate, a larger welfare gain, lower average purchased sugar, and positive reformulation activity. Under the fixed-characteristics benchmark, the optimal rate is $\tau = 0.015$, which implies an 18 cent tax on a 12 g/100ml product. Expressed per 10 grams of sugar, this benchmark is about 2.54 U.S. cents, which is broadly comparable in scale to the 2.28 pence (about 3.1 U.S. cents) fixed-characteristics sugar-tax estimate in O'Connell and Smith's U.K. application. Under our reformulation-allowed optimum, the rate rises to $\tau = 0.040$, or 48 cents on a 12 g/100ml product, while welfare increases from \$119.7 million to \$784.9 million and average purchased sugar falls from 6.27 to 1.74 g/100ml.

Table 2.10: Optimal Tax Choice: Fixed Characteristics vs. Endogenous Reformulation

Scenario	Firms can reform.	τ	Tax on 12g (\$/20oz)	Welfare (\$m)	Avg. sugar (g/100ml)	Reform. units
Fixed characteristics	No	0.015	0.18	119.7	6.27	0
Endog. reformulation	Yes	0.040	0.48	784.9	1.74	88

The tax is τs_j per 20-oz equivalent, where s_j is sugar in grams per 100ml. The first row reports the fixed-characteristics benchmark in the spirit of O'Connell and Smith (2024). The second row reports our model's optimal tax when firms are allowed to reformulate before repricing. *Firms can reform.* indicates whether firms are allowed to adjust sugar content in the counterfactual. Welfare is the present-value change relative to no tax over a 36-month horizon with a 5 percent annual discount rate.

Table 2.11 separates tax choice from the counterfactual used to evaluate the policy. The first row repeats the fixed-characteristics benchmark. The second row keeps the same tax rate, $\tau = 0.015$, but allows firms to reformulate. Comparing these two rows isolates the value of the reformulation margin at a fixed tax rate: welfare rises from \$119.7 million to \$241.5

million and average purchased sugar falls from 6.27 to 5.69 g/100ml. The third row then reports our model's optimal tax when the planner anticipates reformulation, so comparing the second and third rows shows how incorporating reformulation into the policy problem changes both the tax choice and the resulting welfare.

Table 2.11: Tax Choice When Reformulation Is Available

Policy choice	Firms can reform.	τ	Tax on 12g ($/20oz)	Welfare	Exter. savings	Reform. cost	Avg. sugar (g/100ml)	Reform. units
Fixed char. optimum	No	0.015	0.18	119.7	632.0	0.0	6.27	0
Fixed-char. tax	Yes	0.015	0.18	241.5	694.3	56.6	5.69	76
Endog. reform. optimum	Yes	0.040	0.48	784.9	1456.7	355.0	1.74	88

Policy choice names the model used to select the tax rate. *Firms can reform.* indicates whether firms are allowed to adjust sugar content when the policy is simulated. Welfare, externality savings, and reformulation costs are reported in millions of dollars. Welfare is measured relative to no tax and includes consumer surplus, producer surplus, tax revenue, externality savings, and reformulation costs.

Welfare Decomposition

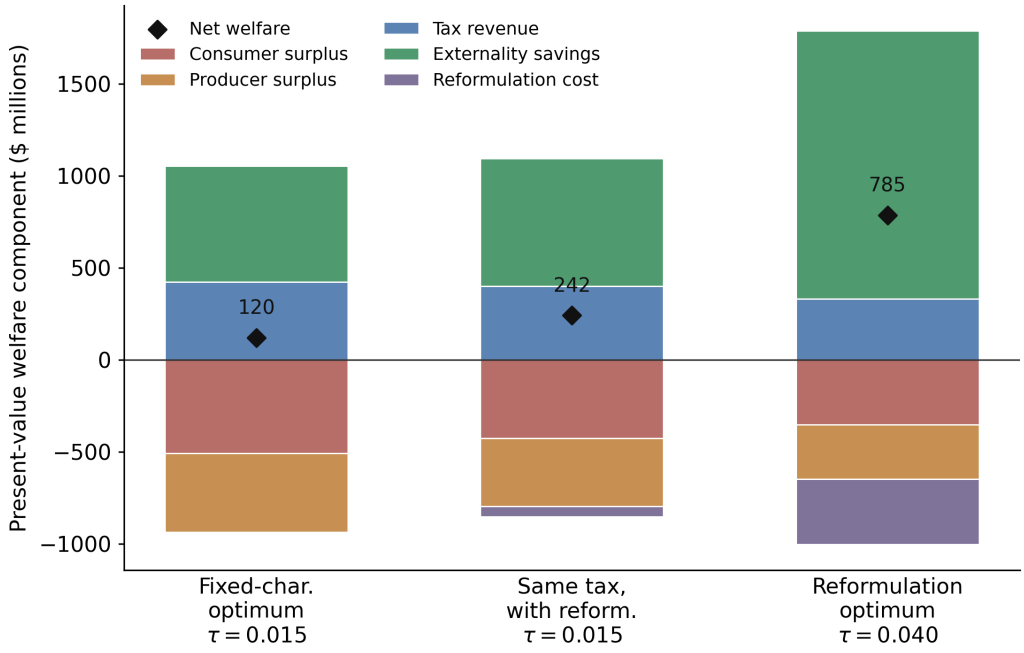
Under our reformulation-allowed optimum,[13] the tax rises to 48 cents on a 12 g/100ml drink and welfare rises to $784.9 million. Figure 2.4 shows where those additional gains come from.

Figure 2.4 helps organize the welfare comparison across the three counterfactuals. Because each bar reflects a different equilibrium outcome, the figure is most informative when read as a comparison of how the welfare components move together across policies. Two comparisons are most useful. First, comparing the first and second bars holds the tax rate fixed at $\tau = 0.015$ and shows how much a fixed-characteristics analysis understates the welfare gain from that tax. Once firms are allowed to reformulate in response to the same tax, net welfare rises from $119.7 million to $241.5 million. The figure shows that the understatement comes mainly from missing externality savings. Consumer and producer surplus are still lower than under no tax, but those losses are smaller once firms can reformulate, while reformulation costs remain modest. Second, comparing the second and third bars shows what happens when the planner reoptimizes the tax once reformulation

[13]The counterfactual uses the central 35 percent reformulation-cost calibration, $\kappa = 0.449$, and holds fixed baseline formulation-cost residuals that rationalize the observed no-tax product formulations, so the no-tax baseline remains an equilibrium object before the tax is introduced. Appendix B.3 gives the simulation details for the optimal-tax counterfactuals.

is anticipated. The higher tax induces much larger externality savings and much higher reformulation costs, while tax revenue falls somewhat. Consumer and producer surplus remain below their no-tax levels, but their losses are smaller than under the lower tax. In the California application, the planner therefore prefers a higher tax not because reformulation is costless, but because the additional reduction in sugar-related harm more than offsets the extra private costs of inducing it.

Figure 2.4: Welfare Decomposition for Linear Sugar Taxes



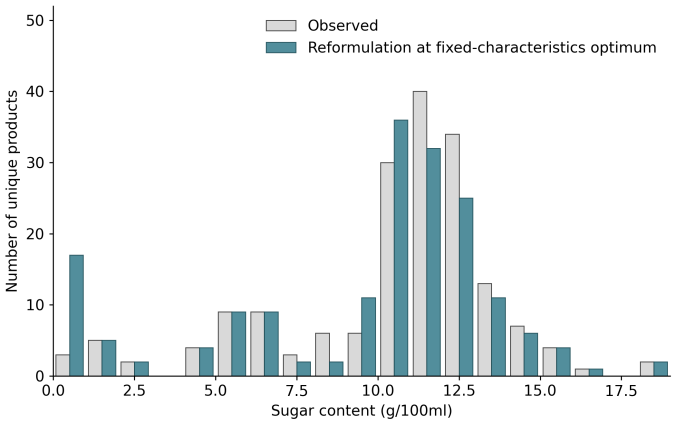
Notes: The figure decomposes present-value welfare changes relative to no tax over the 36-month counterfactual horizon. Externality savings are the negative of the change in external damages. Reformulation costs enter welfare negatively. Diamonds report net welfare changes.

Product Reformulation

The sugar-distribution figures show how reformulation occurs in product space.

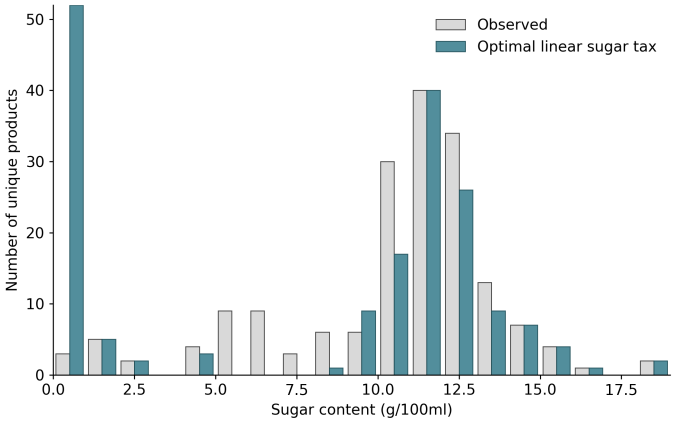
Because the histograms are based on counts of unique products in the reformulation choice set, they are best read as evidence on the breadth and direction of product redesign rather than on sales-weighted importance. Under the tax chosen by the fixed-characteristics model, Figure 2.5 already shows some large downward adjustments: the mass of near-zero-sugar products increases, but the main cluster of products remains centered around 10 to

Figure 2.5: Sugar Content Under the Fixed-Characteristics Optimum



Notes: The figure plots the distribution of sugar content across unique products in the reformulation choice set under observed formulations and under the fixed-characteristics optimum, $\tau = 0.015$. Bars report counts of unique products in one-gram bins.

Figure 2.6: Sugar Content Under the Endogenous-Reformulation Optimum



Notes: The figure plots the distribution of sugar content across unique products in the reformulation choice set under observed formulations and under the endogenous-reformulation optimum, $\tau = 0.040$. Bars report counts of unique products in one-gram bins.

13 g/100ml. Under the reformulation-allowed optimum, Figure 2.6 shows a much stronger response. The number of near-zero-sugar products rises sharply, while product counts fall across a range of mid- and high-sugar bins. The model therefore predicts that reformulation does not occur as a uniform trimming of sugar across all products. Instead, a subset of products makes large reductions in sugar content, and that subset expands when the tax is

set optimally.

Taken together, the decomposition and distribution figures show why reformulation is first-order in the policy problem. A fixed-characteristics analysis misses this margin twice: it understates the welfare gain from a given tax, and it recommends a tax rate that is too low once firms are able to redesign products.

2.9 Conclusion

This paper studies optimal sugar taxation when firms can endogenously reformulate differentiated products. We develop a theory of sin taxation in which taxes affect welfare through both consumer substitution and firms' choice of sugar content, so the optimal policy depends not only on external harms, market power, and redistribution, but also on the cost of reformulation. To quantify that margin, we estimate a structural model of demand and supply for the California non-alcoholic beverage market and use product-level evidence from the U.K. Soft Drinks Industry Levy to calibrate the reformulation-cost parameter, which is not observed in the California data.

The quantitative results show that reformulation is first-order for policy design. In an efficiency benchmark, a fixed-characteristics counterfactual selects an optimal linear sugar tax of 18 cents on a 12 g/100ml drink. Incorporating manufacturers' reformulation responses raises the optimum to 48 cents, increases welfare gains from $119.7 million to $784.9 million over a 36-month horizon, and lowers average purchased sugar from 6.27 to 1.74 g/100ml. Even at the fixed-characteristics tax rate, allowing firms to reformulate raises welfare to $241.5 million, implying that a policy analysis that abstracts from reformulation misses a substantial share of the gains from sugar taxation. Those additional gains come mainly from larger externality savings as the tax induces a broad shift of products toward lower-sugar formulations.

More broadly, the paper suggests that the design of sin taxes should depend on the supply-side margins they create. When reformulation is feasible, taxes on sugar content can be substantially more effective than a fixed-characteristics analysis would imply, and threshold-based designs may be especially attractive when policy aims to induce discrete reductions in sugar content around salient cutoffs.

References

1. Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky. 2019a. "Should We Tax Sugar-Sweetened Beverages? An Overview of Theory and Evidence." Journal of Economic Perspectives, 33(3): 202–227.

2. Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky. 2019b. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." Quarterly Journal of Economics, 134(3): 1557–1626.

3. Barahona, Nano, Cristobal Otero, and Sebastian Otero. 2023. "Equilibrium Effects of Food Labeling Policies." Econometrica, 91(3): 839–868.

4. Barwick, Panle Jia, Hyuk-Soo Kwon, and Shanjun Li. 2024. "Attribute-Based Subsidies and Market Power: An Application to Electric Vehicles." NBER Working Paper 32264.

5. Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica, 63(4): 841–890.

6. Cawley, John, David Frisvold, David Jones, and Chelsea Lensing. 2021. "The Pass-Through of a Tax on Sugar-Sweetened Beverages in Boulder, Colorado." American Journal of Agricultural Economics, 103(3): 987–1005.

7. Cawley, John, David Frisvold, Anna Hill, and David Jones. 2020. "The Impact of the Philadelphia Beverage Tax on Prices and Product Availability." Journal of Policy Analysis and Management, 39(3): 605–628.

8. Cawley, John and David E. Frisvold. 2017. "The Pass-Through of Taxes on Sugar-Sweetened Beverages to Retail Prices: The Case of Berkeley, California." Journal of Policy Analysis and Management, 36(2): 303–326.

9. Dickson, Alex, Markus Gehrsitz, and Jonathan Kemp. 2025. "Does a Spoonful of Sugar Levy Help the Calories Go Down? An Analysis of the UK Soft Drinks Industry Levy." Review of Economics and Statistics, 107(6): 1754–1763.

10. Dubois, Pierre, Rachel Griffith, and Martin O'Connell. 2020. "How Well Targeted Are Soda Taxes?" American Economic Review, 110(11): 3661–3704.

11. Eichner, Thomas, and Marco Runkel. 2025. "Efficient Taxation of Sugar-Sweetened Beverages and Missing Markets for Sugar Content." Scandinavian Journal of Economics, 127(3): 576–612.

12. Eizenberg, Alon. 2014. "Upstream Innovation and Product Variety in the U.S. Home PC Market." Review of Economic Studies, 81(3): 1003–1045.

13. Fan, Ying, and Chenyu Yang. 2020. "Competition, Product Proliferation, and Welfare: A Study of the U.S. Smartphone Market." American Economic Journal: Microeconomics, 12(2): 99–134.

14. Fan, Ying, and Chenyu Yang. 2022. "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety." NBER Working Paper 30146.

15. Grummon, Anna, Benjamin B. Lockwood, Dmitry Taubinsky, and Hunt Allcott. 2019. "Designing Better Sugary Drink Taxes." Science, 365(6457): 989–990.

16. Hendel, Igal, and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." Econometrica, 74(6): 1637–1673.

17. HM Treasury. 2018. "Soft Drinks Industry Levy Comes into Effect." Press release, April 5. London: GOV.UK.

18. Jones-Smith, Jessica C., Melissa A. Knox, Norma B. Coe, Lina P. Walkinshaw, John Schoof, Deven Hamilton, Philip M. Hurvitz, and James Krieger. 2022. "Sweetened Beverage Taxes: Economic Benefits and Costs According to Household Income." Food Policy, 110: 102277.

19. O'Connell, Martin, and Kate Smith. 2024. "Optimal Sin Taxation and Market Power." American Economic Journal: Applied Economics, 16(4): 34–70.

20. Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." Journal of Political Economy, 110(4): 705–729.

21. Rojas, Christian, and Emily Wang. 2021. "Do Taxes on Soda and Sugary Drinks Work? Scanner Data Evidence from Berkeley and Washington State." Economic Inquiry, 59(1): 95–118.

22. Wang, Y. Claire, Pamela Coxson, Yu-Ming Shen, Lee Goldman, and Kirsten Bibbins-Domingo. 2012. "A Penny-Per-Ounce Tax on Sugar-Sweetened Beverages Would Cut Health and Cost Burdens of Diabetes." Health Affairs, 31(1): 199–207.

23. World Health Organization. 2023. *Global report on the use of sugar-sweetened beverage taxes, 2023.* Geneva: WHO.

APPENDIX A

Complementarity and Product Variety: Appendix

A.1 Additional Tables and Figures

Table A.1.1: Statista Consumer Insights Sample for Ownership Moments

Panel A. Survey respondents by wave					
Module	2019	2020	2021	2022	Total
Smartphone brand ownership	3,965	4,013	6,667	9,522	24,167
Tablet brand ownership	2,928	2,813	4,469	6,304	16,514
Dual-owner cross-ownership	2,832	2,736	4,315	6,052	15,935

Panel B. Pooled brand composition (%), 2019–2022			
Module	Apple	Samsung	Other
Smartphone respondents	45.5	29.5	25.0
Tablet respondents	45.3	25.1	29.6

Notes: Source is Statista Consumer Insights, U.S. adults aged 18–64, fielded in annual waves. Panel A reports the survey base of respondents reached by each module in each year; the smartphone module asks for the brand of the primarily used smartphone (single pick), the tablet module asks for tablet brands owned (multi-pick), and the dual-owner cross-ownership module is the subset of respondents who own at least one device of each type. Panel B reports unweighted shares of respondents whose mapped brand is Apple, Samsung, or Other, pooled across 2019–2022. The brand-ownership moments enter the estimation with 41,780 mapped brand records pooled across the smartphone and tablet modules; the difference relative to the sum of bases reflects the multi-pick structure of the tablet module and the inclusion of an "Other" aggregator in the mapped denominator. The cross-ownership moments enter the estimation with 17,565 dual-owner records jointly mapped to the (Apple, Samsung, Others)×(Apple, Samsung, Others) grid, pooled as a separate 2019 wave and a 2020–2022 wave.

Table A.1.2: Full Micro-Moment Fit

Moment	Data	Model
Income share moments		
Phone Apple non-high-price share, Middle–Low income	0.132	0.092
Phone Apple non-high-price share, High–Low income	0.268	0.231
Phone Apple high-price share, Middle–Low income	0.052	0.067
Phone Apple high-price share, High–Low income	0.134	0.155
Phone Samsung non-high-price share, Middle–Low income	0.030	0.038
Phone Samsung non-high-price share, High–Low income	-0.015	0.009
Phone Samsung high-price share, Middle–Low income	-0.035	-0.046
Phone Samsung high-price share, High–Low income	-0.113	-0.118
Tablet Apple non-high-price share, Middle–Low income	0.128	0.120
Tablet Apple non-high-price share, High–Low income	0.251	0.249
Tablet Apple high-price share, Middle–Low income	0.041	0.060
Tablet Apple high-price share, High–Low income	0.065	0.113
Tablet Samsung non-high-price share, Middle–Low income	0.031	0.002
Tablet Samsung non-high-price share, High–Low income	0.008	-0.021
Tablet Samsung high-price share, Middle–Low income	-0.028	-0.054
Tablet Samsung high-price share, High–Low income	-0.051	-0.103
Income paid-price moments		
Phone Apple paid price, Middle–Low income	38.7	52.9
Phone Apple paid price, High–Low income	57.8	95.3
Phone Other paid price, Middle–Low income	62.6	42.5
Phone Other paid price, High–Low income	115.1	77.3
Tablet Apple paid price, Middle–Low income	15.9	49.7
Tablet Apple paid price, High–Low income	41.2	163.2
Tablet Other paid price, Middle–Low income	22.3	16.4
Tablet Other paid price, High–Low income	61.5	25.7
Cross-device ownership moments		
Apple tablet share among Apple-phone owners, 2019	0.656	0.605
Apple tablet share among Apple-phone owners, 2020–2022	0.579	0.647
Samsung tablet share among Samsung-phone owners, 2019	0.371	0.325
Samsung tablet share among Samsung-phone owners, 2020–2022	0.392	0.351
Samsung tablet share among Apple-phone owners, 2019	0.106	0.117
Samsung tablet share among Apple-phone owners, 2020–2022	0.122	0.085
Apple tablet share among Samsung-phone owners, 2019	0.275	0.246
Apple tablet share among Samsung-phone owners, 2020–2022	0.224	0.220

Continued on next page

Table A.1.2: Full Micro-Moment Fit (continued)

Moment	Data	Model
Apple tablet share among other-brand phone owners, 2019	0.248	0.193
Apple tablet share among other-brand phone owners, 2020–2022	0.205	0.191
Samsung tablet share among other-brand phone owners, 2019	0.278	0.192
Samsung tablet share among other-brand phone owners, 2020–2022	0.254	0.167
Brand ownership moments		
Apple smartphone owner share, 2019–2022	0.455	0.394
Samsung smartphone owner share, 2019–2022	0.295	0.245
Apple tablet owner share, 2019–2022	0.453	0.393
Samsung tablet owner share, 2019–2022	0.251	0.175
Ownership-rate moments		
Smartphone ownership rate, 2015	0.690	0.571
Smartphone ownership rate, 2018	0.770	0.624
Smartphone ownership rate, 2021	0.850	0.653
Tablet ownership rate, 2013	0.340	0.431
Tablet ownership rate, 2015	0.450	0.542
Tablet ownership rate, 2018	0.530	0.542
Tablet ownership rate, 2021	0.530	0.542
Replacement moments		
Smartphone replacement cycle	2.71	2.18
Tablet replacement cycle	4.60	4.76

Notes: The table reports all 49 active micro-moments used in demand estimation. Data are empirical moments; Model reports fitted values from the estimated model. Share, ownership, and ownership-rate moments are probabilities. Paid-price differences are in real dollars, and replacement moments are in years. Income share and paid-price moments are differences relative to low-income consumers; non-high-price denotes prices below the product-specific high-price cutoff.

Table A.1.3: OpenBrand Survey Sample for Income Moments

Panel A. Sample size and price cutoff

Model market	OpenBrand purchase category	Years	Priced purchases	Weighted purchases	High-price cutoff
Smartphones	Cell phones	2015–2022	294,225	290,524	$596.59
Tablets	Tablets/e-readers	2015–2022	173,505	168,643	$365.58

Panel B. Weighted composition (%)

Model market	Apple	Samsung	Other brand	Low income	Middle income	High income
Smartphones	32.6	31.3	36.1	48.1	27.7	24.2
Tablets	37.1	21.3	41.6	44.2	28.7	27.0

Notes: The table summarizes the priced OpenBrand respondent-purchase records used to construct income micro-moments. High-price cutoffs are the product-market-specific 80th percentiles of CPI-normalized paid prices used to split purchases into high-price and non-high-price groups. Composition shares use survey weights and sum to 100 within market up to rounding.

Table A.1.4: Highest-Quality Smartphone and Tablet Products, 2022Q4

Rank	Brand	Model	Quality index	Screen	PPI	CPU	RAM
Panel A. Smartphones							
1	Sony	Xperia 1 IV	6.28	6.5	643	3.00	16.0
2	Samsung	Galaxy Z Fold4	6.13	7.6	373	3.20	12.0
3	Samsung	Galaxy S21 Ultra	5.97	6.8	516	2.90	12.0
4	Sony	Xperia PRO-I	5.96	6.5	643	2.80	12.0
5	Sony	Xperia 1 III	5.89	6.5	643	2.80	12.0
6	Sony	Xperia Pro	5.79	6.5	643	2.80	12.0
7	Samsung	Galaxy Z Fold3	5.78	7.6	372	2.80	12.0
8	Samsung	Galaxy S22 Ultra	5.70	6.8	379	3.00	11.0
9	ZTE	Axon 40 Pro	5.61	6.7	393	3.20	10.0
10	Xiaomi	MI 12T Pro	5.56	6.7	444	3.20	12.0
Panel B. Tablets							
1	Samsung	Galaxy Tab S8 Ultra	8.66	14.6	239	3.00	12.0
2	Lenovo	Tab P12 Pro	8.15	12.6	240	3.20	8.0
3	Samsung	Galaxy Tab S8+	8.12	12.4	266	3.00	9.3
4	Samsung	Galaxy Tab S7+	7.98	12.4	266	3.09	7.0
5	Apple	iPad Pro (2022)	7.91	12.9	265	3.49	11.2
6	Apple	iPad Pro (2021)	7.67	12.9	265	3.20	11.2
7	Apple	iPad Pro (2022)	7.66	11.0	265	3.49	11.2
8	Apple	iPad Pro (2021)	7.41	11.0	265	3.20	11.2
9	Samsung	Galaxy Tab S8	6.64	11.0	274	3.00	9.3
10	Samsung	Galaxy Tab S7	6.50	11.0	274	3.09	7.0

Notes: The table reports the ten highest-quality model-level product-quarter observations in the 2022Q4 IDC smartphone and tablet samples. Products are ranked within category by the category-specific PCA quality index used in the demand, marginal-cost, fixed-cost, and product-variety analyses. PPI is pixels per inch, CPU is processor speed in GHz, RAM is measured in GB, and screen size is measured in inches. Repeated commercial names correspond to distinct model-level products in the estimation data.

Table A.1.5: Demand Responses to a $10 Price Increase

Panel A. Smartphones

Price of	13	14 Pro Max	14 Pro	A13	S22	S22 Ultra	Pixel 6a	Moto Power	Moto Stylus
13	-6.46	0.43	0.43	0.21	0.27	0.30	0.26	0.20	0.24
14 Pro Max	0.31	-6.15	0.35	0.12	0.19	0.23	0.16	0.11	0.14
14 Pro	0.27	0.30	-6.32	0.11	0.17	0.20	0.14	0.11	0.13
A13	0.09	0.07	0.07	-8.03	0.18	0.15	0.16	0.18	0.17
S22	0.08	0.08	0.08	0.13	-7.26	0.17	0.11	0.09	0.10
S22 Ultra	0.09	0.09	0.09	0.10	0.16	-6.75	0.09	0.07	0.08
Pixel 6a	0.10	0.09	0.09	0.16	0.14	0.13	-7.86	0.16	0.16
Moto Power	0.03	0.02	0.03	0.06	0.04	0.03	0.06	-8.43	0.06
Moto Stylus	0.03	0.03	0.03	0.06	0.05	0.04	0.05	0.06	-8.16

Panel B. Tablets

Price of	iPad 2021	Air 2022	mini 2021	Tab A8	S7 FE	S6 Lite	HD 10	HD 8	HD 8 Plus
iPad 2021	-5.34	1.33	1.47	1.23	1.27	1.28	1.14	0.64	0.73
Air 2022	0.96	-4.69	1.29	0.55	0.95	0.65	0.35	0.11	0.14
mini 2021	0.29	0.35	-6.07	0.18	0.26	0.21	0.14	0.06	0.07
Tab A8	0.42	0.26	0.33	-9.57	0.94	1.20	1.26	0.66	0.69
S7 FE	0.05	0.06	0.06	0.11	-7.65	0.12	0.06	0.02	0.03
S6 Lite	0.05	0.04	0.04	0.14	0.11	-8.79	0.09	0.05	0.05
HD 10	0.38	0.16	0.24	1.21	0.45	0.73	-10.25	0.93	0.92
HD 8	0.15	0.03	0.07	0.44	0.12	0.27	0.65	-10.11	0.88
HD 8 Plus	0.14	0.04	0.07	0.37	0.12	0.24	0.52	0.72	-9.94

Notes: Cells report the local percent change in the column product's market share from a $10 increase in the row product's price, evaluated in 2022Q3. Rows index the product whose price changes, and columns index the product whose market share responds. Within each device category, displayed products are the top three Apple products, top three Samsung products, and top three non-Apple/non-Samsung products by market share in the evaluation quarter.

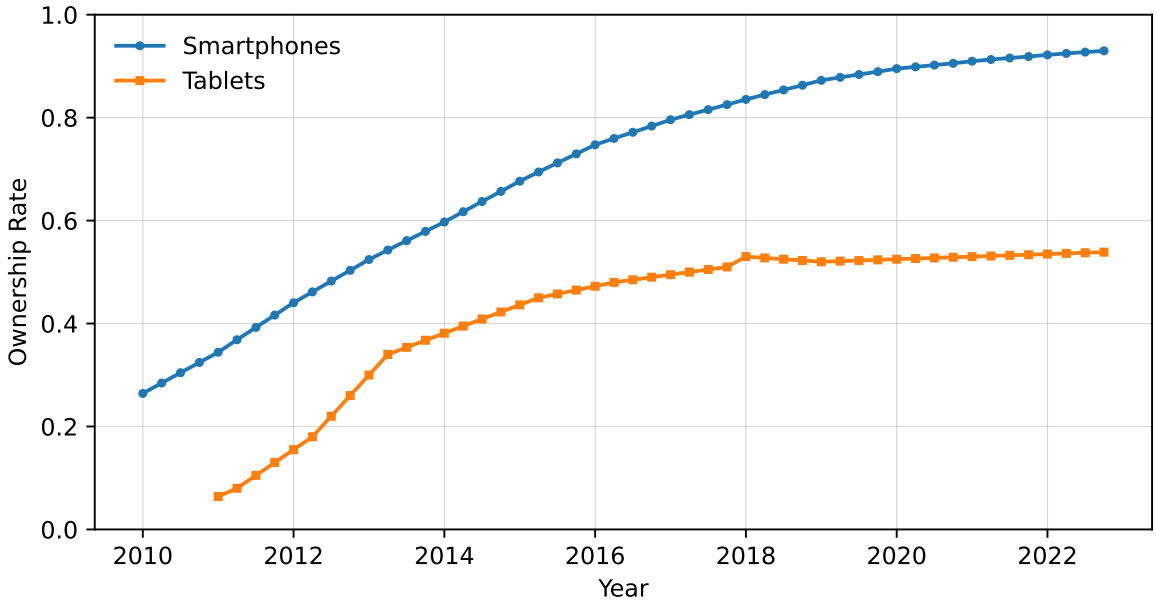
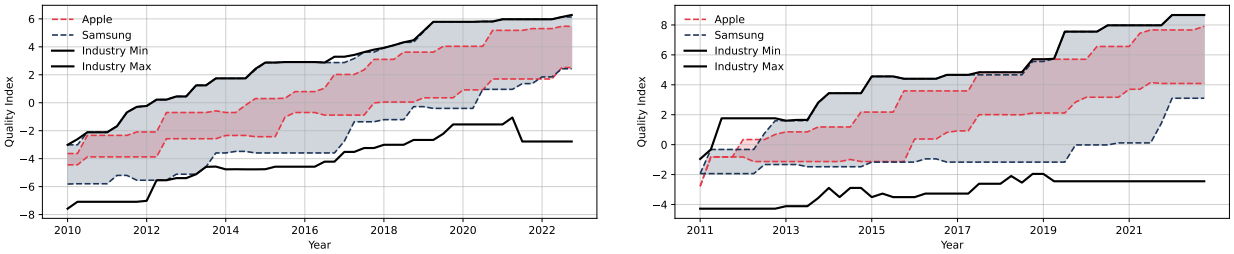


Figure A.1.1: Ownership Rates Used in the Ownership Micro-Moments

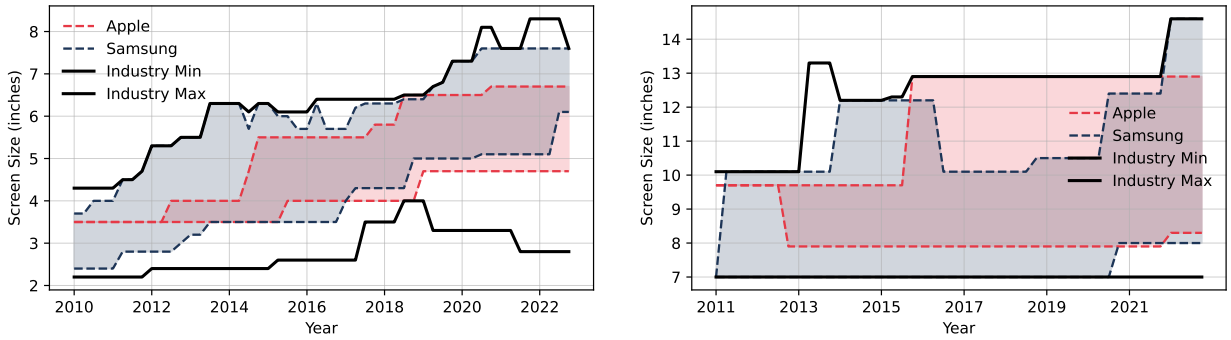
Notes: The figure plots the quarterly smartphone and tablet ownership rates used to construct the Pew ownership-rate moments. Annual survey observations are assigned to product-market quarters through the interpolation used in the estimation data.



(a) Smartphones (b) Tablets

Figure A.1.2: Minimum-to-Maximum Quality-Index Range by Brand

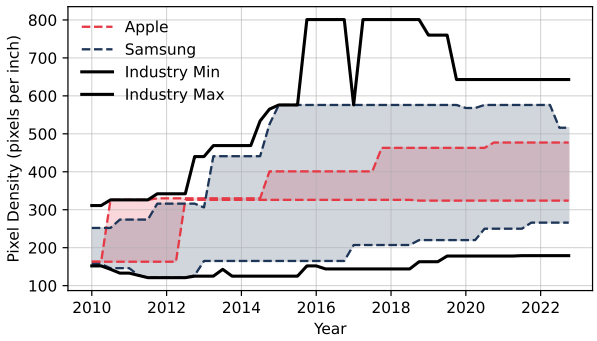
Notes: Shaded regions show the brand-level minimum-to-maximum range of the quality index by quarter for Apple and Samsung. Solid black lines mark the industry-wide minimum and maximum across all firms. The quality index is constructed separately for smartphones and tablets from standardized hardware characteristics.



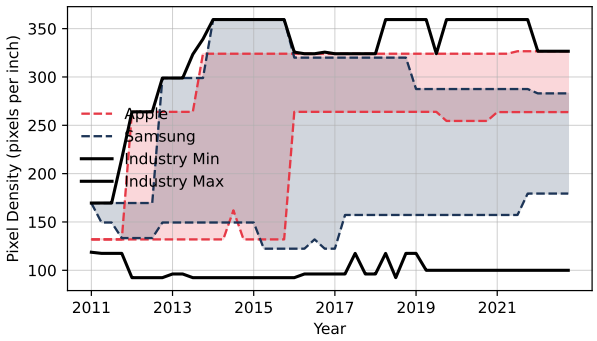
(a) Smartphones (b) Tablets

Figure A.1.3: Range of Screen Sizes by Brand

Notes: Shaded regions show the brand-level minimum-to-maximum range of screen sizes (inches) by quarter for Apple and Samsung. Solid black lines mark the industry-wide minimum and maximum across all firms.



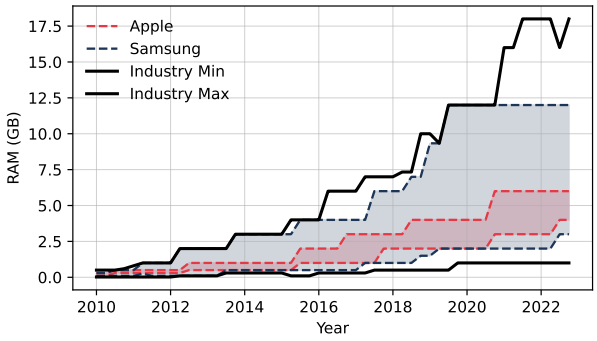
(a) Smartphones



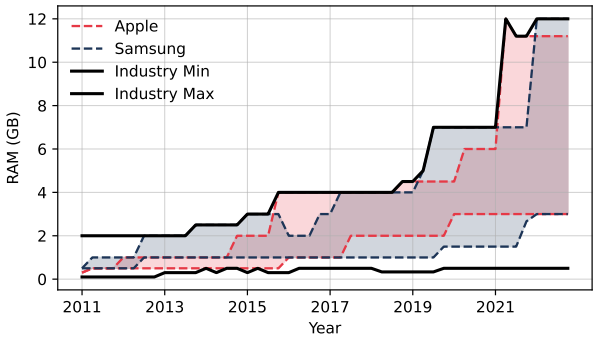
(b) Tablets

Figure A.1.4: Range of Pixel Density by Brand

Notes: Shaded regions show the brand-level minimum-to-maximum range of pixel density by quarter for Apple and Samsung. Solid black lines mark the industry-wide minimum and maximum across all firms. The plotted values are in pixels per inch.



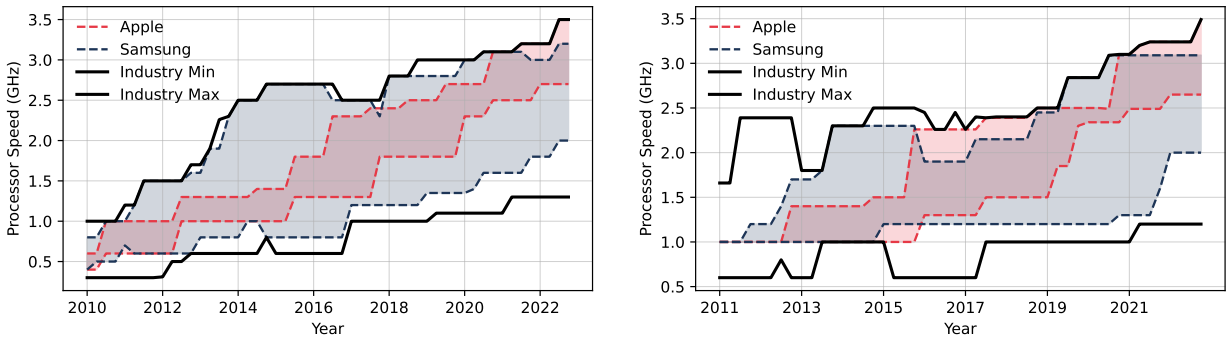
(a) Smartphones



(b) Tablets

Figure A.1.5: Range of RAM by Brand

Notes: Shaded regions show the brand-level minimum-to-maximum range of RAM (GB) by quarter for Apple and Samsung. Solid black lines mark the industry-wide minimum and maximum across all firms.



(a) Smartphones

(b) Tablets

Figure A.1.6: Range of Processor Speed by Brand

Notes: Shaded regions show the brand-level minimum-to-maximum range of processor speed (GHz) by quarter for Apple and Samsung. Solid black lines mark the industry-wide minimum and maximum across all firms.

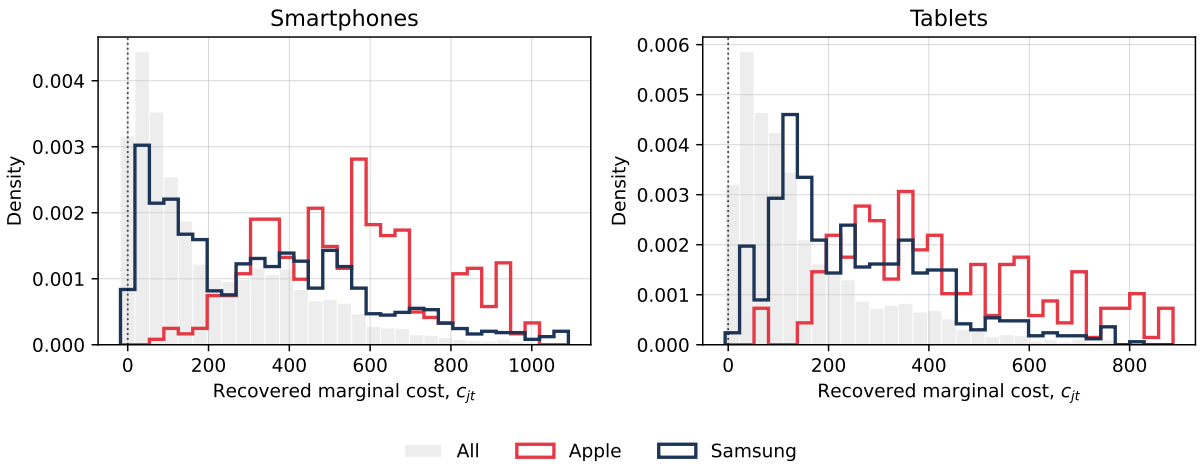


Figure A.1.7: Distribution of Recovered Marginal Costs

Notes: The figure plots density histograms of recovered product-level marginal costs, c_{jt}, using product-quarter observations in the final estimation sample. Gray bars use all products in the category; red and blue outlines overlay Apple and Samsung products. The dotted vertical line marks zero. Histogram ranges trim the lower and upper 0.5 percent of the plotted variable for readability; negative-cost shares reported in the text are computed on the full estimation sample.

A.2 Complementarity Feature Coding

The complementarity data record feature availability separately from cross-device compatibility. For each product-quarter and feature f, an availability indicator records whether public release notes and support documentation imply that the product can use feature f in that quarter. A separate compatibility group records the ecosystem over which the feature works. A smartphone–tablet pair receives feature f in the complementarity index only when both devices have the feature and their compatibility groups match. Pairs involving no previously owned device receive zero feature complementarity.

Cloud Sync maps iCloud and Samsung Cloud. Apple cloud support begins in 2011Q4 for eligible iPhones and iPads, following the release of iOS 5 and the contemporaneous iCloud rollout. Samsung cloud support enters in 2017Q1 for supported Galaxy phones and tablets. The compatibility groups distinguish Apple iCloud, Samsung Cloud, and products without supported cloud synchronization.

Share maps cross-device file sharing. Apple sharing begins in 2013Q4 for AirDrop-eligible iPhones and iPads. Samsung and Android sharing enters later, beginning around the 2020 rollout of Quick Share and Nearby Share for supported Galaxy and Android devices. The sharing compatibility group is Apple for AirDrop and Android-based for supported non-iOS devices.

Continuity maps handoff, call forwarding, and text-message continuity. Apple continuity begins in 2014Q4 for iOS 8-supported devices. Samsung continuity begins in 2015Q2 for eligible Samsung phones and tablets. The compatibility groups distinguish Apple Continuity, Samsung continuity, and unsupported devices.

Potential products in the fixed-cost and counterfactual exercises keep these rules when eligibility can be inferred from brand, model family, operating system, release timing, and product category. For newly drawn model-quality alternatives without product-specific support pages, eligibility is assigned on the external quality scale. Apple cloud support is available after the iCloud rollout, and Apple phone sharing and continuity use the quality thresholds implied by the iPhone 4S and iPhone 4 eligibility rules. Samsung and other non-Apple support is assigned when the new product's quality exceeds the minimum quality of an in-data supported product for the same firm-market and feature. Otherwise, unsupported potential-product features are coded as zero. These rules make potential-product compatibility depend on local empirical support rather than on unrestricted feature adoption.

A.3 Micro-Moment Definitions

This appendix defines the micro-moment analogues used in estimation.

Let w_i denote simulation weights, and let $\mathcal{D}_{bt}^m$ denote products of type $m \in \{P, T\}$ in brand group b that can be owned in period t, with $\mathcal{D}_t^m = \cup_b \mathcal{D}_{bt}^m$. The brand-level ownership probability is the sum of the product-level ownership probabilities defined in Section 1.4.1.1, $\sum_{d \in \mathcal{D}_{bt}^m} p_{idt}^m(\theta)$.

For phone brand group $b \in \{\text{Apple}, \text{Samsung}, \text{Others}\}$ and tablet brand group $b' \in \{\text{Apple}, \text{Samsung}\}$, the Statista conditional ownership moment is

$$\begin{aligned} &\Pr_\theta(\text{own tablet brand } b' \mid \text{own phone brand } b, \text{own a tablet}; \tau) \\ &\quad = \frac{\sum_{t \in \mathcal{T}_\tau} \sum_i w_i \left(\sum_{d \in \mathcal{D}_{bt}^P} p_{idt}^P(\theta)\right) \left(\sum_{d' \in \mathcal{D}_{b't}^T} p_{id't}^T(\theta)\right)}{\sum_{t \in \mathcal{T}_\tau} \sum_i w_i \left(\sum_{d \in \mathcal{D}_{bt}^P} p_{idt}^P(\theta)\right) \left(\sum_{d' \in \mathcal{D}_t^T} p_{id't}^T(\theta)\right)}. \end{aligned} \tag{A.3.1}$$

The Statista marginal brand-ownership moment is

$$\Pr_\theta(\text{own brand } b \mid \text{own type } m; \tau) = \frac{\sum_{t \in \mathcal{T}_\tau} \sum_i w_i \sum_{d \in \mathcal{D}_{bt}^m} p_{idt}^m(\theta)}{\sum_{t \in \mathcal{T}_\tau} \sum_i w_i \sum_{d \in \mathcal{D}_t^m} p_{idt}^m(\theta)}, \qquad b \in \{\text{Apple}, \text{Samsung}\}. \tag{A.3.2}$$

The Pew ownership-rate moment is

$$\Pr_\theta(\text{own type } m; \tau) = \frac{1}{|\mathcal{T}_\tau|} \sum_{t \in \mathcal{T}_\tau} \sum_i w_i \sum_{d \in \mathcal{D}_t^m} p_{idt}^m(\theta). \tag{A.3.3}$$

Let $a_{i1}^m < a_{i2}^m < \cdots < a_{iK_i^m}^m$ denote the simulated acquisition periods for consumer i in category m. The replacement-cycle moment is the average time, in years, between consecutive acquisitions:

$$E_\theta(\text{years between acquisitions of type } m) = \frac{\sum_i w_i \mathbf{1}\{K_i^m \geq 2\} \left[\frac{1}{K_i^m - 1} \sum_{k=2}^{K_i^m} \frac{a_{ik}^m - a_{i,k-1}^m}{4}\right]}{\sum_i w_i \mathbf{1}\{K_i^m \geq 2\}}. \tag{A.3.4}$$

For the income moments, let m_i denote the product type purchased by consumer i, b_i the purchased brand, y_i income, and p_i the real paid price. The income bins are $\mathcal{Y}^{Low}$, $\mathcal{Y}^{Middle}$, and $\mathcal{Y}^{High}$. For each $m \in \{P, T\}$, $\mathcal{P}_m^{High}$ collects purchases with p_i above the 80th percentile of real paid prices, and $\mathcal{P}_m^{Other}$ collects the remaining purchases. For $\mathcal{Y} \in \{\mathcal{Y}^{Middle}, \mathcal{Y}^{High}\}$,

$\mathcal{P} \in \{\mathcal{P}_m^{Other}, \mathcal{P}_m^{High}\}$, and $b \in \{\text{Apple}, \text{Samsung}\}$, the brand-share moment is

$$\begin{aligned} &\Pr_\theta(\text{brand} = b \mid \text{purchase type} = m, p_i \in \mathcal{P}, y_i \in \mathcal{Y}) \\ &\quad - \Pr_\theta(\text{brand} = b \mid \text{purchase type} = m, p_i \in \mathcal{P}, y_i \in \mathcal{Y}^{Low}). \end{aligned} \tag{A.3.5}$$

The brand-conditional paid-price moment is

$$\begin{aligned} &E_\theta[p_i \mid \text{purchase type} = m, \text{brand} = b, y_i \in \mathcal{Y}] \\ &\quad - E_\theta[p_i \mid \text{purchase type} = m, \text{brand} = b, y_i \in \mathcal{Y}^{Low}], \qquad b \in \{\text{Apple}, \text{Other}\}, \end{aligned} \tag{A.3.6}$$

where Other pools all non-Apple and non-Samsung brands.

A.4 Ownership Simulation Details and Monte Carlo Validation

This appendix documents the ownership-state simulation used in demand estimation and the Monte Carlo design behind Table 1.4. The Monte Carlo exercise tests whether aggregate shares and simulated ownership histories recover a complementarity parameter when individual ownership transitions are unobserved. The design is simpler than the full empirical model, but keeps the relevant state-dependence problem: two durable-goods categories, evolving ownership states, persistent quality tastes, and a complementarity shock that turns on mid-sample.

A.4.1 Empirical Ownership Simulation Settings

The empirical demand estimation uses 500 base simulated consumers and eight ownership-state paths for each base consumer, for a total of 4,000 simulated rows per product category and market. The base consumers are common across the smartphone and tablet panels. A scrambled Halton sequence with seed 999 generates the persistent taste nodes and the market-specific ownership-transition draws. The persistent taste draws are held fixed for each base consumer and repeated across the eight state paths, while the transition draws vary across paths and markets. Base simulation weights are normalized to sum to one and then divided equally across the eight paths.

Quarterly real income is attached before the ownership paths are expanded. Annual nominal household income follows a persistent latent-rank process with persistence 0.95 and seed 271828. Each year's simulated income cross-section is rank-matched to the CPS ASEC HINC-06 distribution, with uniform draws inside closed income bins and a Pareto tail in the top open bin. Nominal income is fixed within a calendar year and deflated quarter by quarter using the CPI. The demand specification uses standardized log real income.

Ownership states are initialized to no owned product before each category enters the sample; in the simulation arrays, product id -1 denotes no owned device. In the first smartphone market, the ownership probabilities are rescaled to match the observed ownership rate. After that initial calibration, inside purchases replace the consumer's previous device in that category and outside-option draws leave the previous device in place. These simulated ownership histories generate both the ownership probabilities used in the moment analogues and the realized ownership paths used for replacement moments.

A.4.2 Monte Carlo Design

The data-generating process uses the empirical smartphone and tablet product-market skeleton, holding products, prices, observed characteristics, and market timing fixed. The three empirical compatibility features are replaced by a single scalar complementarity index, equal to zero before 2016Q3 and active thereafter. Simulated demand follows a random-coefficients logit with price, quality, brand, age-bin controls, product-market demand shocks, a scalar complementarity parameter, and a random coefficient on quality. Simulated consumers begin with no owned device in either category; each inside purchase replaces the previously owned device in that category, while an outside-option draw leaves ownership unchanged.

The estimator observes simulated aggregate product shares and product-market characteristics, but not individual ownership transition matrices. For each candidate parameter vector, it simulates ownership histories, solves market-by-market intercepts to match aggregate inside shares, and predicts product-level market shares. It then estimates the standard demand parameters, the scalar complementarity parameter, and the quality-taste dispersion by minimizing product-level log-share residuals.

The exercise uses 10 replications for each complementarity scale, 0, 0.5, 1, and 1.5, where the scale multiplies the baseline scalar complementarity parameter of 1. Each replication uses 500 simulated consumers, 9,987 smartphone product-market observations, 3,863 tablet product-market observations, and 52 quarterly markets, of which 48 include tablets. The true quality-taste dispersion is $\sigma_q = 0.5$, and the product-market demand shock standard deviation is 0.2. Aggregate shares are computed from model choice probabilities rather than finite-sample purchase counts.

Table A.4.1: Monte Carlo Recovery of Standard Demand Parameters

Scale	Price	Quality	Apple	Samsung	Age 2	Age 3
True	-1.000	0.600	0.400	0.200	-0.250	-0.500
0.0	-1.006	0.603	0.405	0.206	-0.254	-0.505
0.5	-1.007	0.605	0.406	0.202	-0.253	-0.500
1.0	-1.005	0.602	0.400	0.203	-0.253	-0.506
1.5	-1.006	0.604	0.401	0.204	-0.254	-0.503

Notes: The table reports mean estimates across 10 replications for each complementarity scale. The estimator jointly recovers these mean-utility parameters, the scalar complementarity parameter, and σ_q.

The additional mean-utility recovery estimates reported in Table A.4.1 show that the

standard mean-utility parameters are centered close to their true values across complementarity scales. Together with Table 1.4, these results indicate that the estimator recovers complementarity without attributing the same variation to quality-taste dispersion or standard product characteristics in this controlled environment.

The Monte Carlo validates the ownership-state mapping under the maintained transition rule. It does not test robustness to misspecified replacement, churn, depreciation, initial ownership distributions, or equilibrium supply responses to unobserved demand shocks.

A.5 Fixed-Cost and Counterfactual Details

This appendix defines the support and simulation objects used in the fixed-cost and portfolio counterfactual exercises: revealed-preference intervals, the potential product set, admissible fixed-cost residuals, and the counterfactual equilibrium computation.

A.5.1 Fixed-Cost Bounds and Moment Inequalities

Observed portfolios imply revealed-preference bounds through one-product add and drop deviations. Drop deviations for products offered in the data generate upper bounds on fixed costs; add deviations for feasible absent products generate lower bounds. The payoff comparison is the finite-horizon expected profit statistic from the supply model, so the bounds incorporate both own-market variable profits and cross-market spillovers through future ownership states.

The bound sample contains 10,187 smartphone product-market intervals and 5,071 tablet product-market intervals before instrument-bin stacking. Direct inequalities supply 4,904 smartphone lower endpoints, 5,283 smartphone upper endpoints, 3,379 tablet lower endpoints, and 1,692 tablet upper endpoints. Raw negative bounds are set to zero, and all bounds are measured in millions of real dollars, before support completion.

The support-completion step fills the missing side of each one-sided interval within cells defined by year, firm bracket, within-year quality bin, Quality Distance bin, and model-age group. Apple and Samsung retain separate firm brackets, while all other firms are pooled. Model age is grouped as 0–3, 4–7, and 8 or more quarters. Within each cell, a missing lower endpoint is filled with the minimum finite direct bound in the cell and a missing upper endpoint is filled with the maximum finite direct bound in the cell. These completion cells are used only to form finite product-market intervals; the moment-instrument bins below are separate nonnegative instrument functions.

Let X_{jt}^{FC} denote the fixed-cost covariates and let γ^m denote the fixed-cost coefficient vector for product category $m \in \{P, T\}$. For a completed interval $[\underline{FC}_{jt}, \overline{FC}_{jt}]$ and nonnegative instrument function $z_g(X_{jt}^{FC})$, the unconditional moment inequalities are

$$\mathbb{E}\Big[z_g(X_{jt}^{FC})\left(X_{jt}^{FC\prime}\gamma^m - \underline{FC}_{jt}\right)\Big] \geq 0, \qquad \mathbb{E}\Big[z_g(X_{jt}^{FC})\left(\overline{FC}_{jt} - X_{jt}^{FC\prime}\gamma^m\right)\Big] \geq 0.$$

The empirical moments use bins of model quality, Quality Distance, model age, and brand group. Tablet moments also include quality-by-age bins. The stacked moment system has

40,748 lower and 40,748 upper smartphone rows across 14 instrument groups, and 25,355 lower and 25,355 upper tablet rows across 26 instrument groups.

The fixed-cost index includes a constant, the model-level quality index, Quality Distance, Apple and Samsung indicators, and indicators for model ages 4–7 quarters and 8 or more quarters. Quality Distance is the absolute distance between the product's quality index and the nearest product already offered by the same firm in that product market. The omitted model-age group is 0–3 quarters. All fixed-cost objects are measured in millions of real dollars per model-quarter.

A.5.2 Potential Product Set

The potential set $\mathcal{K}_{nt}^{m}$ is defined at the model level. It combines recently discontinued models with newly drawn model-quality alternatives. This support matches the product-choice margin in the supply model: firms choose device models, while storage and RAM configurations are absorbed into model-quarter characteristics.

Recently discontinued models remain in the potential set for two quarters after their final observed appearance. These carry-forward products keep their model-level quality index, age forward, and remain attached to the firm that offered them. They capture continuation of a model just outside the observed portfolio.

New model-quality alternatives are drawn from local empirical support. For a firm-market, the support is

$$\left[\max\{q_{mt}^{\min}, q_{nt}^{\min} - 2\}, \min\{q_{mt}^{\max}, q_{nt}^{\max} + 1\}\right],$$

where $q_{mt}^{\min}$ and $q_{mt}^{\max}$ are the market-level quality endpoints and $q_{nt}^{\min}$ and $q_{nt}^{\max}$ are the firm's endpoints in that market. A draw is accepted only if it is at least 0.2 quality-index units from every model already offered by the firm in the market. The fixed-cost bound exercise uses eight accepted draws per firm-market and doubles that number for Apple and Samsung. The terminal-market counterfactual support uses the same rule with five accepted draws per firm-market, again doubled for Apple and Samsung.

Newly drawn products enter as observed model-level products do. They enter with age zero, inherit the firm's major-brand fixed-effect category when available, and receive imputed initial prices from observed price-quality-age relationships with common year and quarter effects. Imputed prices are floored at \$30 for smartphones and \$50 for tablets. Quality Distance is recomputed after the potential set is formed, so add-deviation inequalities use

the distance from each candidate to the firm's contemporaneous portfolio. Figure A.5.1 illustrates this draw rule for the 2022Q4 market.

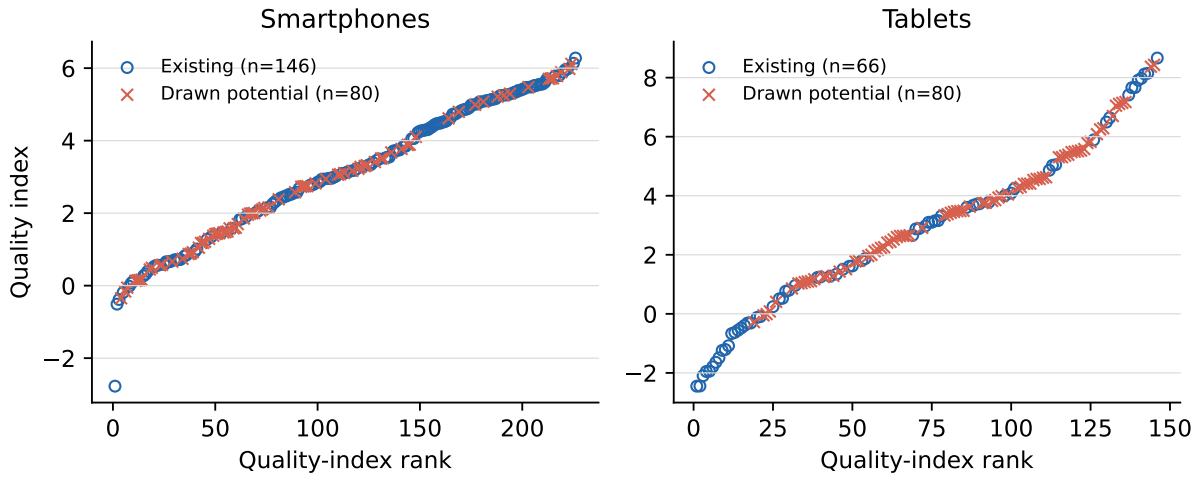


Figure A.5.1: Example Potential-Product Quality Draws

Notes: The figure uses 2022Q4 as an example market. Circles are observed model-level products and crosses are accepted newly drawn potential products. The horizontal axis ranks observed and newly drawn products together by the quality index within each product category. The figure isolates the new-quality draw rule; the full potential set also includes recently discontinued carry-forward products.

A.5.3 Fixed-Cost Residual Calibration

The portfolio counterfactual implementation constructs the systematic fixed-cost component using the componentwise midpoints of the projected identified intervals for γ^m. For each product j in the terminal-market potential set, write fixed cost as

$$FC_j = X_j^{FC\prime}\gamma^m_{\text{mid}} + \nu_j,$$

where γ^m_{mid} denotes the midpoint vector and ν_j is a product-level residual. The residuals are drawn to satisfy the baseline revealed-preference inequalities, so the observed 2022Q3 portfolio remains consistent with no profitable one-product deviations before the counterfactual policy is applied.

For each residual draw, expected variable profits are computed for the one-product deviations in the portfolio game. These deviations translate the baseline no-profitable-deviation conditions into residual bounds. Observed products impose drop-deviation upper bounds;

feasible absent products impose entry-deviation lower bounds. Because a direct one-product inequality usually identifies only one side of a product's fixed cost, the support is completed before drawing residuals.

This terminal-market residual completion is distinct from the support completion used to estimate the fixed-cost inequalities. The direct residual bounds are converted into total fixed-cost bounds, and missing endpoints are filled using the empirical minimum and maximum of finite direct add/drop total fixed-cost bounds in the same product category, firm bracket, and quality-bin cell, where the firm bracket treats major firms separately and pools other firms. If this primary cell produces a zero-width interval, the completion uses the broader product-category by quality-bin cell.

The completed interval $[\underline{FC}_j, \overline{FC}_j]$ is then converted back into a residual interval

$$\left[\underline{FC}_j - X_j^{FC\prime}\gamma^m_{\text{mid}},\, \overline{FC}_j - X_j^{FC\prime}\gamma^m_{\text{mid}}\right].$$

Residuals are drawn independently across products from uniform distributions on these product-specific intervals. The draws are therefore independent but not identically distributed. Infeasible completed intervals stop the counterfactual run. The reported counterfactuals average over four admissible fixed-cost residual draws, and the same residual vectors are used across the baseline and policy environments.

A.5.4 Counterfactual Simulation Settings

The reported portfolio counterfactuals evaluate the 2022Q3 market for ASUS, Amazon, Apple, LG, Motorola/Lenovo, and Samsung. Each policy changes the complementarity environment and holds fixed the estimated demand, marginal-cost, and fixed-cost primitives, the ownership-state simulation, and the product universe. The endogenous-portfolio equilibrium recomputes both product offerings and prices.

Firms update assortments through the sequential one-product best-response algorithm of Fan and Yang (2020). In each firm-category step, the firm compares the current portfolio with all feasible single-product additions and removals in that category, taking rivals' portfolios in both categories as given. The payoff is expected profit across smartphones and tablets, averaged over a 12-quarter horizon with equal period weights. Each expected-profit calculation uses 10 antithetic demand and marginal-cost shock draws.

The best-response routine cycles through firms and categories until no one-product deviation raises expected profit, up to 50 rounds. Within a firm-category best response, the

routine can accept up to 20 successive one-product improvements before moving to the next firm-category pair. Conditional on the final portfolio, prices are resolved and outcomes are evaluated using common fresh demand and marginal-cost shocks across baseline and policy environments. Reported outcomes include consumer surplus, profits net of fixed costs, fixed costs, product counts, and prices.

APPENDIX B

Optimal Sin Tax with Endogenous Reformulation: Appendix

B.1 Numerical Mechanism Exercise

This appendix describes the normalized four-product duopoly used in the numerical mechanism exercise. The purpose of the exercise is to clarify the conditions under which a notched tax can outperform a linear sugar tax; it is not intended as a quantitative calibration to the California application. Each firm owns two products, and each product has a baseline sin characteristic s_j^0, a mean quality δ_j, a constant marginal cost mc_j, and a reformulation-cost scale.

Table B.1.1: Toy Simulation Primitives

Product	Firm	δ_j	s_j^0	mc_j	Cost scale
A_classic	Firm A	0.78	12.0	0.20	0.0130
A_light	Firm A	0.64	8.0	0.17	0.0095
B_classic	Firm B	0.74	13.0	0.21	0.0145
B_fruit	Firm B	0.58	10.0	0.18	0.0115

Notes: Sugar s_j is measured in grams per 100ml. Cost scale is the product-specific reformulation-cost parameter used below. The market size is normalized to one.

Demand is logit. Product j's deterministic utility is

$$u_j(p_j, s_j) = \delta_j - \gamma_s(s_j - s_j^0)^2 - \alpha p_j, \qquad \alpha = 2.4, \quad \gamma_s = 0.0075. \tag{B.1.1}$$

The quadratic term captures the idea that moving away from the baseline formulation

can reduce product appeal. Marginal cost is fixed at mc_j in the central exercise. This isolates reformulation incentives from mechanical marginal-cost changes. External harm is proportional to aggregate sin consumption, $0.10 \sum_j s_j q_j$.

For each policy, firms reoptimize prices and product characteristics. The stage-2 pricing equilibrium is the multiproduct Bertrand equilibrium under the relevant taxes and sugar choices. Welfare is

$$W = CS + PS + TR - 0.10 \sum_j s_j q_j, \tag{B.1.2}$$

where producer surplus is variable profit net of any reformulation costs.

The exercise compares two tax instruments. The linear sugar tax sets $\tau_j = \tau s_j$ and searches over $\tau \in [0, 0.12]$. The single-notch volumetric tax sets $\tau_j = r\mathbf{1}\{s_j > \bar{s}\}$ and searches over $r \in [0, 1.20]$ and a threshold grid that includes zero, baseline sugar levels, policy-relevant sugar levels, and midpoints. The central run used 41 linear-tax grid points, 11 rate grid points for r, and 13 threshold points. The grid is designed to compare mechanisms rather than to estimate calibrated policy magnitudes.

The supply side is solved under three reformulation environments. With fixed characteristics, firms only reprice and $s_j = s_j^0$. With zero-cost continuous reformulation, firms choose $s_j \in [0, s_j^0]$ without reformulation costs. This case serves as an upper-bound diagnostic. With smooth costs, product j pays

$$RC_j^{smooth}(s_j, s_j^0) = k_j \left(\frac{s_j - s_j^0}{4} \right)^2. \tag{B.1.3}$$

For the linear tax, the continuous-cost models solve firms' best responses over sugar and prices.

For the notched tax, the discontinuity makes local first-order conditions unreliable at the threshold. The toy therefore evaluates policy-implied candidate formulations: each product can remain at baseline, move to the threshold from above, or move to zero. This candidate set captures the bunching margin that the notched tax is designed to create under the normalized primitives. For each candidate threshold and rate, the code enumerates pure equilibria and records the welfare-maximizing selected equilibrium. This is the object reported in Table 2.1.

The numerical checks verify the internal consistency of the supply models and their shared zero-tax allocation. In the central run, the maximum price first-order-condition residual is below 2×10^{-13}, maximum profitable-deviation gains are zero, and the zero-tax

comparability check has zero welfare and sugar gaps.

B.2 Reformulation-Cost Calibration Details

The calibration applies the original U.K. Soft Drinks Industry Levy schedule to the California products. For a product with sugar content s_j in grams per 100ml, the levy per litre is

$$\tau^{SDIL}(s_j) = \begin{cases} 0, & s_j < 5, \\ 0.18 \text{ pounds/litre}, & 5 \leq s_j < 8, \\ 0.24 \text{ pounds/litre}, & s_j \geq 8. \end{cases} \tag{B.2.1}$$

The model's quantity unit is a 20-oz equivalent. The benchmark conversion uses 0.591 litres per 20-oz equivalent and an exchange rate of 1.335 dollars per pound, so the model tax is

$$tax_j^{SDIL} = \tau^{SDIL}(s_j) \times 0.591 \times 1.335. \tag{B.2.2}$$

The candidate set reflects the two tax notches in the SDIL schedule:

$$S_j = \begin{cases} \{s_j\}, & s_j \leq 5, \\ \{s_j, 4.9\}, & 5 < s_j < 8, \\ \{s_j, 7.9, 4.9\}, & s_j \geq 8. \end{cases} \tag{B.2.3}$$

The values 4.9 and 7.9 are computational proxies for moving just below the relevant threshold. Products below the lower threshold are not used to match the reformulation-rate target because they are not initially levy-eligible.

For each eligible product j and candidate $s' \in S_j$, the calibration changes only product j's sugar content and leaves all other products' sugar contents fixed. It then applies the SDIL schedule, updates marginal cost for product j, and solves the stage-2 pricing equilibrium. With the linear marginal-cost specification in the main text, the marginal-cost update is

$$\Delta mc_j = \eta^{sugar}(s'_j - s_j) + \eta^{sugar \times juice} x_j^{juice}(s'_j - s_j) + \eta^0 \left[\mathbf{1}\{s'_j = 0\} - \mathbf{1}\{s_j = 0\} \right]. \tag{B.2.4}$$

The SDIL candidates 4.9 and 7.9 do not change the zero-sugar indicator, but the same formula is used for consistency with the estimated supply equation.

The benchmark profit object is firm-level rather than own-product profit. For each market where formulation j is sold, the profit gain sums the change in variable profit over every

product owned by the same firm. This captures within-firm cannibalization when product j reformulates.

The benchmark computes profit gains in the last observed month across all California FIPS and treats that month as a representative monthly profit environment. The one-month profit gain is scaled to a 36-month present value using a 5 percent annual discount rate:

$$PV\Delta\Pi_j(s') = \Delta\Pi_j^{month}(s') \sum_{m=1}^{36} \frac{1}{(1+r_m)^m}, \qquad r_m = (1.05)^{1/12} - 1. \tag{B.2.5}$$

The same present-value factor scales the baseline quantity Q_j in the reformulation-cost term. This keeps the profit benefit and the quantity-scaled adjustment cost on the same payoff horizon.

For a candidate κ, each eligible product chooses the sugar level with the highest net value:

$$V_j(s';\kappa) = PV\Delta\Pi_j(s') - \kappa Q_j \left(\frac{s' - s_j}{10}\right)^2. \tag{B.2.6}$$

The product reformulates if the best candidate differs from s_j. Exact ties are assigned to no reformulation. The calibrated κ is the grid value that makes the count-weighted reformulation share closest to the target:

$$\hat{\kappa} = \arg\min_{\kappa} \left| \frac{1}{|\mathcal{E}|} \sum_{j \in \mathcal{E}} \mathbf{1}\{\hat{s}_j(\kappa) \neq s_j\} - R^{target} \right|, \qquad \mathcal{E} = \{j : s_j > 5\}. \tag{B.2.7}$$

Sales-weighted reformulation shares are reported as diagnostics, but the benchmark matches the count-weighted target because the external evidence is closest to a product-level reformulation moment.

The benchmark uses 159 eligible reformulation units observed in 54 county markets in the last month of 2018. All baseline pricing equilibria and all 303 candidate-level pricing equilibria converged under the pricing solver. The maximum attainable count-weighted reformulation share at $\kappa = 0$ is 75.5 percent; the corresponding sales-weighted share is 36.4 percent. No eligible products are dropped by the within-unit sugar-stability filter in the benchmark run.

B.3 Optimal Linear Sugar-Tax Simulation Details

The optimal-tax counterfactuals use the estimated demand and marginal-cost primitives from the structural model and recompute equilibrium prices under each tax. The current benchmark uses the last observed month across all California FIPS as a representative monthly market environment, giving 54 markets and 40,650 product–market rows. Products are measured in 20-oz equivalents, and sugar is measured in grams per 100ml. For a linear sugar tax τ, product j in market m pays

$$tax_{jm}(\tau, s_j) = \tau s_j \tag{B.3.1}$$

per 20-oz equivalent.

The counterfactual first normalizes the no-tax baseline to be a pricing equilibrium. For each market, observed prices and observed sugar contents imply predicted shares and a multiproduct Bertrand markup vector $\eta_m(p^0, s^0)$. Baseline marginal costs for the counterfactual are set to

$$c_m^0 = p_m^0 - \eta_m(p_m^0, s_m^0). \tag{B.3.2}$$

This normalization leaves the estimated demand system unchanged but makes the observed no-tax prices a fixed point of the same pricing map used in the policy simulations. When a product changes sugar content, marginal cost changes according to the estimated sugar-cost terms:

$$c_{jm}(s_j) = c_{jm}^0 + \eta^{sugar}(s_j - s_j^0) + \eta^{sugar\times juice} x_j^{juice}(s_j - s_j^0) + \eta^0 \left[\mathbf{1}\{s_j = 0\} - \mathbf{1}\{s_j^0 = 0\}\right]. \tag{B.3.3}$$

Given a sugar vector s, the stage-2 equilibrium price vector solves the multiproduct Bertrand fixed point

$$p_m = c_m(s) + \tau s_m + \eta_m(p_m, s_m), \tag{B.3.4}$$

where $\eta_m(p_m, s_m)$ is computed from the random-coefficients demand derivatives and the observed firm ownership matrix. The demand system is evaluated at counterfactual prices and sugar contents relative to the estimated baseline utilities. The mean-utility and heterogeneous-utility components are updated for price, sugar, sugar squared, and the estimated price and sugar heterogeneity terms.

The fixed-characteristics benchmark holds $s_j = s_j^0$ for every product. For each candidate τ, the simulation solves (B.3.4) in every market, computes shares, profits, tax revenue,

external damages, and consumer surplus, and then selects the tax rate with the highest present-value welfare over the grid. There is no reformulation cost in this benchmark because product characteristics are fixed.

The endogenous-reformulation benchmark adds a stage-1 sugar-content game before pricing. The reformulation choice set contains products whose observed sugar content is stable within the counterfactual sample, whose baseline sugar is positive, and whose observations are assigned to a single firm. In the current all-FIPS run this gives 179 reformulation units. A reformulation unit may choose any sugar content in $[0, s_j^0]$; sugar increases are not allowed. Products outside the reformulation choice set remain in the market and enter demand, pricing, and firm profits, but their sugar contents are held fixed.

For each tax rate, firms play a sequential best-response algorithm over their eligible reformulation units. Given rival sugar choices, firm f chooses the vector s_f for all eligible units in its portfolio to maximize

$$A\Pi_f(s_f, s_{-f}; \tau) - RC_f(s_f, s_f^0), \qquad A = \sum_{h=1}^{36} \frac{1}{(1+r_m)^h}, \quad r_m = (1.05)^{1/12} - 1. \tag{B.3.5}$$

The profit term is the firm's total variable profit across all of its products after re-solving the stage-2 pricing equilibrium. This is the margin that allows the firm to internalize cannibalization between a reformulated product and other products in its portfolio. The reformulation cost combines the calibrated smooth cost with the one-sided baseline formulation residual:

$$RC_f(s_f, s_f^0) = \sum_{j \in \mathcal{J}_f^R} \left[\kappa Q_j A \left(\frac{s_j - s_j^0}{10} \right)^2 + \eta_j Q_j A \max\{s_j^0 - s_j, 0\} \right]. \tag{B.3.6}$$

The smooth component uses the 35 percent calibration, $\kappa = 0.449$. The residual η_j is a one-sided formulation-cost term recovered from no-tax deviations. It is the minimum product-level linear cost, within the tested deviation set and additional firm-level cuts, needed to make observed no-tax formulations stable against profitable sugar reductions. These residuals are held fixed in all tax counterfactuals.

Each firm best response is solved as a bounded continuous optimization problem over its eligible sugar vector. A best response is accepted when it increases the firm's present-value payoff by more than the payoff tolerance. The stage-1 game is treated as converged when no accepted best response changes any firm's sugar vector by more than the sugar tolerance

and the maximum remaining payoff gain is below the payoff tolerance. The reported endogenous optimum is selected only among tax points for which all market-level price equilibria converge and the stage-1 reformulation game satisfies this convergence criterion. Negative marginal costs induced by the baseline marginal-cost anchoring are reported as diagnostics, but they are not used as a hard filter in the benchmark optimum selection.

For both supply environments, welfare is computed relative to the no-tax observed baseline. In the efficiency benchmark reported in the main text, $g_i = 1$, $\lambda = 1$, and the profit-tax wedge is set to zero. Monthly welfare for a tax rate τ is

$$\Delta W(\tau) = \Delta CS(\tau) + \Delta\Pi(\tau) + TR(\tau) - \Delta D(\tau), \tag{B.3.7}$$

where $D(\tau)$ is external damage from physical sugar consumption. Damages are 0.007 dollars per gram of sugar, and physical grams equal $s_j \times 5.9147 \times q_{jm}$ for quantity q_{jm} measured in 20-oz equivalents. Present-value welfare is $A\Delta W(\tau) - RC(\tau)$. Consumer surplus is computed from the log-sum change for each simulated consumer draw, using the draw-specific price coefficient to convert utility into dollars and the social weight for that draw. Producer surplus is total variable profit, tax revenue is $\sum_{jm} \tau s_j q_{jm}$, and reformulation costs are paid only in the endogenous-reformulation benchmark.

The fixed-characteristics grid is refined around its welfare peak. The endogenous grid is refined around tax rates where the stage-1 game produces high welfare and satisfies the convergence criterion. The main-text tables report the best converged grid point in each environment rather than fitting a smooth welfare curve through the discrete tax evaluations.