

DYNAMIC DEMAND AND PRICING STRATEGY IN E-BOOK MARKET

Hui Li

A DISSERTATION

in

Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2015

Supervisor of Dissertation

Holger Sieg, The J.M. Cohen Term Professor of Economics

Graduate Group Chairperson

George Mailath, Walter H. Annenberg Professor in the Social Sciences

Dissertation Committee

Eric T. Bradlow, K. P. Chao Professor of Marketing, Statistics, and Education

Ulrich Doraszelski, Professor of Business Economics and Public Policy

Jean-François Houde, Assistant Professor of Business Economics and Public Policy

Xun Tang, Assistant Professor of Economics

DYNAMIC DEMAND AND PRICING STRATEGY IN E-BOOK MARKET

© COPYRIGHT

2015

Hui Li

This work is licensed under the
Creative Commons Attribution
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/3.0/>

ACKNOWLEDGEMENT

First and foremost, I would like to express my sincerest gratitude to my advisor Holger Sieg for his invaluable encouragement, support and guidance. He has greatly helped my development as a researcher from the first day that I took his class. I would like to thank him for his guidance at various stages of my dissertation and constant encouragement during my job market process. I also thank him for always believing in me and encouraging me to pursue my goals.

I am also grateful to other members of my dissertation committee. I have benefited tremendously from my interactions with them. I would like to thank Professor Eric T. Bradlow for introducing me to the field of marketing. I thank him for his detailed comments on my dissertation and his constant support during my job market process. I have learnt from him both as a researcher and as a teacher. I would like to thank Professor Jean-François Houde for consistent guidance and detailed suggestions at various stages of my dissertation. I thank him for being always available and helpful. He taught me a lot in his class and during our meetings. I would like to thank Professor Ulrich Doraszelski for his valuable comments and help. The dissertation was developed based on my proposal for his course in my second Ph.D. year. I am grateful for his suggestions and patience in helping me polish my first research paper. I would like to thank Professor Xun Tang for his helpful comments on my dissertation and during my job market process.

Moreover, I would like to thank Professors Yakov Bart, Ron Berman, Peter Fader, Hanming Fang, Cecilia Fieler, Joseph Harrington, Raghuram Iyengar, Barbara E. Kahn, Robert Meyer, Jagmohan Raju, David Reibstein, Katja Seim, Qiaowei Shen, Andrew Shephard, Michael Sinkinson, Petra Todd, Christophe Van den Bulte, Kenneth I. Wolpin, Pinar Yildirim, Z. John Zhang for their comments and suggestions on my dissertation and their help during my job market process. I would like to thank

all my teachers during my doctoral studies at Penn. I benefited a lot from the training and the vision they provided. I thank Wharton Research Data Services (WRDS) for providing access to the data. I also thank my dear fellow students for their help and company.

Finally, I would especially like to thank my parents, Chuxiang Li and Yunfang Dai, and my husband Jiaqi Xu for their constant encouragement and unconditional love.

ABSTRACT

DYNAMIC DEMAND AND PRICING STRATEGY IN E-BOOK MARKET

Hui Li

Holger Sieg

E-reading has experienced rapid growth in the past few years and has raised new questions. On the supply side, retailers such as Amazon jointly sell e-readers and e-books. It remains unclear how they can coordinate the two products to conduct intertemporal price discrimination (IPD). On the demand side, it remains unclear how much of e-book sales come from cannibalizing print books and how much serve as market expansion to the book business.

I empirically address these questions using individual-level data from 2008 to 2012. I estimate a dynamic structural model of consumer e-reader adoption and subsequent book purchases, including quantity, reading format (e-book or print book), and retailer choices (Amazon, other online retailers, or offline bookstores) in a number of book genres. The estimation reveals two consumer types, avid readers and general readers, who self-select into buying e-readers based on their unobserved heterogeneous book tastes. Compared with general readers, avid readers buy more books, adopt e-readers earlier, and have larger cannibalization rates. The two types also have different relative demand elasticities between e-readers and e-books.

Given the estimated demand system, I simulate the optimal dynamic pricing strategies of e-readers and e-books for the monopolist retailer Amazon who faces forward-looking consumers. I find that Amazon should harvest on e-readers and invest in e-books. Complementarity provides the firm a novel dimension of consumer heterogeneity (the relative demand elasticities between e-readers and e-books) to exploit. The joint IPD strategy provides a better screening device for more profitable consumers and limits consumer's ability to intertemporally arbitrage.

To evaluate the impact of e-books on print book sales, I simulate the world without e-books and compare it with the observed one. I find that 42% of e-book sales come from cannibalizing print book sales and that 58% come from market expansion. Of the cannibalization effect, offline bookstores bear 53% of the cannibalization loss, while Amazon bears 32% and other online retailers bear 15%. I further explore how the impact of e-books would change under alternative pricing arrangements. Overall, the results have managerial implications to publishers, book retailers, and policymakers in the e-book market.

TABLE OF CONTENTS

ACKNOWLEDGEMENT	iii
ABSTRACT	v
LIST OF TABLES	ix
LIST OF ILLUSTRATIONS	x
Chapter 1 : Introduction	1
Chapter 2 : Data and Industry Background	11
2.1 The U.S. E-Book Industry	11
2.2 Data Description	12
2.3 Observed Pricing and Consumption Patterns	15
Chapter 3 : Model Setup	18
3.1 A Simple Two-Period Model	18
3.2 Consumer Problem	20
3.3 Firm Problem	31
Chapter 4 : Demand Estimation and Supply Simulation Method	35
4.1 Likelihood Function	35
4.2 Identification	37
4.3 Estimation and Computational Methods	40
Chapter 5 : Demand-Side Estimation Results	42
5.1 Model Fit	42
5.2 Parameter Estimates	43

Chapter 6 : Supply-Side Policy Simulations	49
6.1 IPD on the Primary Product Only	50
6.2 IPD on Both Products	51
6.3 Discussion	56
Chapter 7 : Cannibalization or Market Expansion? The Impact of E-Books on Print Book Sales	59
7.1 Effect Sizes Under Observed Pricing	60
7.2 Effect Sizes Under Alternative E-Book Prices	63
Chapter 8 : Conclusion	66
APPENDIX	71
BIBLIOGRAPHY	78

LIST OF TABLES

TABLE 1 :	Optimal Book Quantity Solutions	24
TABLE 2 :	Model Fit: Cumulative Kindle Sales/Total Market Size (%) . . .	43
TABLE 3 :	Parameter Estimates	45
TABLE 4 :	Consumer Heterogeneity in Kindle and Book Purchases	48
TABLE 5 :	IPD on Both Products vs. IPD on the Primary Product Only . .	53
TABLE 6 :	Effects of Raising E-book Prices	65

LIST OF ILLUSTRATIONS

FIGURE 1 :	Observed Kindle and E-Book Prices	16
FIGURE 2 :	Book Price by Genre and Format	16
FIGURE 3 :	Consumer Decision Timing	21
FIGURE 4 :	Model Fit: Books	44
FIGURE 5 :	Policy Functions: Investing vs. Harvesting	52
FIGURE 6 :	Fractions of Avid Readers in Kindle Sales	54
FIGURE 7 :	Kindle Sales Over Time	55
FIGURE 8 :	Observed Prices vs. Simulated Prices	57
FIGURE 9 :	Cannibalization and Market Expansion	63
FIGURE 10 :	Computation Algorithm	77

CHAPTER 1

Introduction

There has been significant growth in e-reading since Amazon, the largest on-line book retailer, launched its first e-reading device, the Kindle e-reader, in 2007. Consumers first buy an e-reader and then buy e-books from the same retailer. By providing e-readers that are affordable and pricing e-books of new releases and *New York Times* best sellers at \$9.99, Amazon substantially promoted the diffusion of e-reading. E-book sales enjoyed triple-digit annual growth rates from 2008 to 2011 and accounted for 23.8% of the total trade-book unit sales in 2013 (Book Industry Study Group, BookStats 2011, 2014). Yet print book sales declined during the same period, making publishers concerned that e-book sales grew at the expense of the traditional print book business.¹

The diffusion of e-reading raises questions for retailers and publishers. On the supply side, it remains unclear how retailers can dynamically coordinate the pricing of e-readers and e-books to better conduct intertemporal price discrimination (IPD). The pricing question is not unique to the e-book market; firms in other industries also jointly sell durable primary hardware and complementary software, especially in digital and online businesses (e.g., consoles and video games, Apple TVs and digital

¹See <http://www.engadget.com/2011/04/15/e-book-sales-triple-year-over-year-paper-books-decline-in-every/>.

content in iTunes, razors and blades, printers and cartridges). Complementarity is reflected by the fact that (1) consumers need to buy the primary hardware to consume the complementary software, and (2) the usage intensity of the complementary software drives the adoption of the primary hardware.

Two incentives exist for both hardware and software pricing. For hardware, as in the classic razor-and-blades strategy, firms can set low hardware prices to “invest” and penetrate the market so that they can earn from subsequent software sales. Firms can also exploit consumer heterogeneity and “harvest” on the hardware by conducting IPD; they can open with high prices to skim high-valuation consumers and then cut prices later to appeal to low-valuation consumers. For software, firms have incentives to “invest” in new consumers and “harvest” on existing consumers. As the consumer mix evolves over time, it is potentially beneficial for firms to dynamically price software as well. Furthermore, hardware and software pricing are linked, as software price affects the attractiveness of hardware and the hardware price affects the number of software users. In practice, firms either conduct IPD separately for hardware and software without fully exploiting the link between them (e.g., consoles and video games) or conduct IPD only for hardware, keeping the software price stable (e.g., Amazon Kindle and e-books). The possibility of joint IPD on both hardware and software remains understudied by both researchers and practitioners.

On the demand side, it remains unclear how e-books have reshaped the publishing industry. In particular, how much of e-book sales come from cannibalizing print book sales, and how much would not have occurred and serve as market expansion to the book business? How are different print book retailers affected by the introduction of e-books? In practice, publishers were concerned about cannibalization from low-priced e-books and conspired with Apple to raise e-book prices in 2010, which drew scrutiny from the Department of Justice and caused a broad debate over e-book

pricing.² Evaluating the degrees of cannibalization and market expansion helps us understand how publishers and book retailers such as Amazon and offline bookstores are affected by the introduction of e-books, as well as how they should optimally respond.

This dissertation empirically addresses these questions. It provides insights on the e-book market from both the supply and the demand perspectives. I start with estimating a dynamic demand system of books and e-readers and fit the model to individual transaction data from 2008 to 2012. Given the estimated demand system, I numerically solve for the optimal joint IPD strategies and conduct counterfactuals on cannibalization and market expansion. Firms sell hardware (i.e., e-readers) and software (i.e., e-books) to heterogeneous consumers and jointly conduct dynamic pricing on both products. Forward-looking consumers may anticipate future price changes and intertemporally arbitrage. Firms need to account for this strategic behavior in addition to the harvesting and investing incentives.

In the discrete-continuous demand model, consumers choose whether to buy a new e-reader or upgrade to the latest generation. Consumers then maximize their direct utility from books by choosing book quantity, reading format (e-books or print books), and retailer for print books (Amazon.com, other online retailers, or offline bookstores) in a number of book genres. Instead of using a standard discrete choice model and capturing the book usage in a reduced-form way, I explicitly model book consumption as a continuous choice. Consumers can buy multiple books. Their book usage is endogenized to be a function of their unobserved heterogeneous reading tastes and book prices. Book usage further drives e-reader adoption so that e-book prices affect e-reader attractiveness. Consumers respond to price changes by adjusting the number of books to buy, and whether and when to adopt e-readers. The demand system is estimated without assuming that the observed prices are optimal.

²See <http://www.justice.gov/atr/cases/apple/apple-direct.html>.

I find that consumers are highly heterogeneous in book and e-reader consumption. The data identify two major consumer types: avid readers who have higher unobserved reading tastes and general readers who have lower unobserved reading tastes. Estimation results imply that avid readers buy more books and adopt e-readers earlier. They are less price elastic to both e-readers and e-books than general readers in the absolute terms. Yet avid readers are relatively more price elastic to e-books than to e-readers, while general readers are relatively more price elastic to e-readers than to e-books. The intuition is that avid readers buy more books and spend more on books than on e-readers relative to general readers. They care more about subsequent book prices when buying e-readers. Note that the difference in avid readers' and general readers' relative demand elasticity between the two products is not an imposed assumption; it comes from consumers' endogenous choices in the model. This difference is the key driving force of the supply-side pricing policy.

Given the demand estimates, I numerically solve for the optimal joint IPD strategies and address three questions. What is the optimal e-reader and e-book price trajectory combination? How can firms benefit from joint IPD on both products compared with IPD only on e-readers? When is dynamic pricing better and when is fixed pricing better for e-books? Rather than trying to rationalize the observed pricing strategy, I take a normative view and focus on the scenario in which a monopolist Amazon makes dynamic pricing decisions to maximize total profits from e-readers, print books, and e-books. Both consumers and Amazon are forward-looking in the pure-strategy Markov-perfect Nash equilibrium (MPNE). I focus on the pricing problem and take cost, product availability, and quality from the data. Observed e-reader costs have dropped over time and might drive down prices. I define harvesting (investing) as decreasing (increasing) mark-ups over time to avoid the confusion. Prices can still drop in an investing strategy if costs drop faster than prices.

I find that the optimal joint IPD strategy is to harvest on e-readers and invest in

e-books. If the firm conducts IPD only on e-readers and optimally commits to a fixed e-book price, it should invest in e-readers. Simulation results show that conducting IPD on both products benefits the firm in two ways. First, it offers the firm a better screening device to induce a higher fraction of more profitable consumers to buy. Second, it limits consumers' ability to intertemporally arbitrage by providing incentives both to delay purchase and to buy earlier. The profitability of the joint IPD policy, compared with IPD only on e-readers, depends on the composition of heterogeneous consumers in the initial market. The joint IPD policy increases five-year total profits by \$837 million if 30% of consumers are avid readers. If the fraction of avid readers is too low, it is better to commit to a fixed e-book price.

The novelty of the complementary setting is that the firm can exploit a new dimension of consumer heterogeneity, namely, the relative demand elasticities between the two products. Traditional single-product IPD exploits the heterogeneity in the demand elasticities *across* consumer types: Avid readers are less price elastic than general readers to both e-readers and e-books. The joint IPD policy further exploits the heterogeneity in the relative demand elasticities between the two products *within* each consumer type: Avid readers are more price elastic to e-books than to e-readers, while general readers are more price elastic to e-readers than to e-books. Firms can exploit the two-dimensional heterogeneity to price discriminate. The optimal pricing policy is a function of the penetration rates of consumers in each type. I find that for any given general (avid) reader penetration rate, as the avid (general) reader penetration rate increases, the firm should harvest (invest) on e-readers and invest (harvest) in e-books.³ The overall policy depends on the market composition. In general, the retailer can use different joint price trajectories (i.e., harvesting on e-readers and investing in e-books, or investing in e-readers and harvesting on e-books) to induce different consumer types to purchase. The joint IPD policy serves as a

³Notice that consumer types are still unobserved to the firm. This result is a characteristic of the policy functions, which are functions of the penetration rates in each consumer type.

better screening device for more profitable consumers.

The framework can be applied to other industries in which the usage intensity of the software drives the adoption of the hardware. In other words, consumers are self-selected into buying the hardware given their heterogeneous tastes on the software. This feature is central to the dynamic pricing strategies. Different usage intensity leads to different relative demand elasticities between the two products, which in turn drives the pricing policy. Two modeling decisions are important to get unbiased estimates: (1) accounting for consumer heterogeneity in reading tastes, as different types of consumers generate different book revenues and respond differently to e-book prices, and (2) accounting for the dynamic device adoption decision, as it allows for self-selection based on heterogeneous consumer tastes. Modeling self-selection produces more accurate estimates, as the adoption behavior contains information about the consumer's heterogeneous taste. Meanwhile, ignoring self-selection would yield incorrect model predictions. To see this, consider two pricing arrangements in a static setting: 1) high e-reader prices and low e-book prices, and 2) low e-reader prices and high e-book prices. A model without self-selection ignores the fact that consumers differ in their consumption of the two products and react to the two pricing arrangements differently. The model would thus predict that the two pricing arrangements (at some price levels) induce the same number of homogeneous adopters and make no difference to the retailer. However, my model with self-selection would predict that the former induces more avid-reader adopters, while the latter induces more general-reader adopters. The compositions of adopters under the two scenarios have different pricing implications for the firm. Similar logic applies to the intertemporal case, once I replace "high price" with harvesting and "low price" with investing. Ignoring self-selection would provide biased estimates of the mix of consumers and in turn biased predictions on dynamic pricing strategies and cannibalization and market expansion effects.

In Chapter 7, “Cannibalization or Market Expansion? The Impact of E-Books on Print Book Sales,” I use the demand estimates to evaluate how e-books have reshaped the publishing industry, i.e., the interaction of digital and traditional products in the publishing industry. I first calculate the degrees of cannibalization and market expansion given observed e-reader and e-book prices. I simulate a world without e-books and compare it with the observed one. First, I find that 42% of e-book sales come from cannibalizing print book sales and that 58% come from market expansion. Of the e-books sales that come from cannibalization, 53% would have occurred in offline bookstores, 32% on Amazon, and 15% on other online retailers. Amazon bears a higher cannibalization burden over time as print book sales shift from other retailers to Amazon. Accounting for all the retailers, the gain from market expansion is larger than the loss from cannibalization for the publishers. Second, I find that the effect magnitudes differ across consumer types. Avid readers have higher cannibalization rates than general readers do; a higher percentage of avid readers’ e-book consumption comes from cannibalization. The difference leads to a declining industrywide cannibalization rate over time as more general readers adopt e-readers and start to buy e-books. Although both avid readers and general readers buy more books after adopting e-readers, they do not always bring more profits. The industry benefits more from converting general readers to e-book readers, while Amazon benefits more from converting avid readers. Third, I find that the effect magnitudes differ across genres. “Casual” e-books are stronger substitutes for paperbacks than “lifestyle” and “practical” e-books.

I further ask how the impact of e-books would change under counterfactual pricing arrangements. In particular, I raise e-book prices by \$2, similar to what publishers and Apple did in 2010, and explore whether the publishing industry would benefit from this price change. Fully evaluating the impact of raising e-book prices requires accounting for e-reader pricing response. I solve for the optimal dynamic e-reader

pricing strategies of the monopolist retailer Amazon given different e-book prices. I then compare the market outcomes given new e-book and e-reader prices. The results suggest that increasing e-book prices by \$2 makes the industry worse off. The loss from a smaller market expansion effect outweighs the gain from a smaller cannibalization effect. The key is that Kindle owners buy fewer e-books given higher e-book prices, which lowers the gain for Amazon to convert nonowners to owners. Amazon thus increases Kindle prices, which substantially discourages general readers from adopting Kindles and creating market expansion.

The dissertation contributes to the pricing literature by studying IPD in the complementary product case. There have been studies on dynamic pricing of a single product (e.g., Stokey, 1979, 1981; Besanko and Winston, 1990; Nair, 2007; Hendel and Nevo, 2013; and Lazarev, 2013) and static pricing of complementary products (e.g., Gil and Hartmann, 2009). Little is known, however, about dynamic pricing of complementary products. In a complementary product setting, Leung (1997) and Koh (2006) theoretically study durable product IPD in the existence of a flat-rate complementary product. They show that “investing” incentives can outweigh “harvesting” incentives so that increasing prices of the durable product over time is optimal. Nair (2007) and Liu (2010) empirically study IPD in the video game and console industry. They focus on single-product IPD and abstract from either software pricing or hardware pricing. I diverge from the extant literature by modeling the dynamic pricing decisions of both products. In particular, I also allow for IPD on the complementary product and for self-selection based on heterogeneous tastes. This enables me to discover new joint IPD strategies that take advantage of the complementarity.

I also contribute to the literature on complementary products, including tying and bundling. Most studies in this area are developed in a static setting (e.g., Gil and Hartmann 2009). Recent studies empirically explore dynamic demand of complementary products (e.g., Hartmann and Nair, 2010; Sriram, Chintagunta, and Agarwal,

2010; and Derdenger and Kumar, 2013). This dissertation extends the literature by modeling both the dynamic demand and supply pricing problem. It also contributes to the nascent empirical literature on dynamic pricing problems in which both firms and consumers are forward-looking (e.g., Nair, 2007; Goettler and Gordon, 2011). State-of-the-art numerical methods allow me to solve for a model in which consumers' dynamic e-reader adoption is endogenous to the retailer's pricing strategy rather than an exogenously evolving diffusion process. The demand model shares features with dynamic models of technology adoption (e.g., Gowrisankaran and Rysman, 2012; Lee, 2013; and Melnikov, 2013).

The counterfactual on the impact of e-books contributes to the literature on the interactions between the Internet and brick-and-mortar economies. There has been evidence of cannibalization between online newspapers and physical ones (Gentzkow, 2007), YouTube viewing and television viewing (Waldfogel, 2007), file sharing and record sales (Oberholzer-Gee and Strumpf, 2007), and PDF and print format (Kannan, Pope, and Jain, 2009). I complement this literature by empirically examining the e-book case and taking into account the device adoption decision.

This dissertation provides managerial implications for publishers, Amazon, and other print book retailers in the book market. For publishers, the overall market expansion effect outweighs the cannibalization effect. I find that both avid readers and general readers read more after they become e-reader owners based on the model estimates. E-books still make the industry better off, despite that this format causes a redistribution of sales among retailers. For Amazon, I propose a novel joint IPD strategy to better coordinate the dynamic pricing of e-books and e-readers. I also provide some explanation as to why Amazon prices Kindles and e-books at or below cost. Although Amazon bears the cannibalization loss as a print book retailer, it benefits from e-books in three ways: (1) additional e-book sales because of market expansion, (2) additional print book sales driven by e-reader adopters (i.e., e-books

accelerate the shift of book sales from other retailers to Amazon), and (3) spillover effect to other product categories. If gains are large enough, Amazon would have strong incentives to promote e-book and e-reader sales, even at a loss. For policy-makers, they should be aware that publishers and Amazon might have different, even conflicting, incentives and strategies for developing e-reading. They should also take a long-term view about the impact of e-books on the publishing industry.

CHAPTER 2

Data and Industry Background

2.1. The U.S. E-Book Industry

The e-book market did not experience rapid growth until Amazon released its first e-reader, the Kindle, in 2007. Since then, the market size of e-books has grown from \$20 million to \$969.9 million in 2011 (Association of American Publishers, 2012). Amazon's existing relationship with publishers enables it to offer a wide variety of e-books. By providing affordable e-readers at higher qualities over time and pricing e-books of new releases and *New York Times* best sellers at \$9.99, Amazon's market share reached nearly 90% by the end of 2009. Barnes & Noble entered the e-book market in October 2009, accounting for about 20% of all e-book sales by 2011, and has struggled to remain profitable. Apple started to sell e-books in iBookstore and accounted for only about 10% of total e-book sales (Gilbert, 2014). A survey on consumer e-reading shows that the Kindle is still the dominant device used (Bowker Market Research, 2012). As Kindle enjoyed a monopoly position from 2007 to 2009 and was the dominant e-reader from 2008 to 2012, I focus on Amazon's optimal IPD strategies as a monopoly of e-readers and e-books. Consumers can still buy print books from all major retailers: Amazon.com, other online retailers, and offline bookstores.

There are three vertical players (publishers, retailers, and consumers) and three relevant products (print books, e-books, and e-readers) in the market. Publishers sell books to retailers at a wholesale price, which remained stable during the sample period. Retailers then set book retail prices to consumers. Retailers also launched their own e-readers and set e-reader prices.⁴ Discussion with industry practitioners suggests that print book launching and pricing are unaffected by e-book pricing. Thus, I take print book prices as exogenously given from the data. Publishers' wholesale prices are also taken as given. As a robustness check, I allow publishers to optimally change wholesale prices. The predictions on dynamic pricing for retailers remain qualitatively the same.

2.2. Data Description

I combine three individual-level online transaction data sets and supplement them with data on aggregate offline book sales, cost, and e-book availability. The first data set is the individual-level online book transaction records from 2008 to 2012 gathered by comScore.⁵ Each purchase record contains the retail website, purchase

⁴Kindle prices are always determined by Amazon, while two pricing contracts exist for books. Under the wholesale contract, Amazon sets book retail prices and pays wholesale prices to the publishers. Under the agency contract, publishers set book retail prices and Amazon obtains 30% of the book revenue. Print book pricing always follows the wholesale contract. For e-books, there has been a debate over which contract to adopt. Amazon started off with the wholesale contract and set low e-book prices from 2007 to 2010. Publishers were concerned about the cannibalization effect of low-priced e-books. They forced Amazon to sign the agency contract from 2010 to 2012 and raised e-book prices of the new releases and *New York Times* best sellers from \$9.99 to \$12.99-\$14.99. This drew scrutiny from the Department of Justice. The contract scheme was switched back to the wholesale model by the Department of Justice after the lawsuit in 2012 (<http://www.nysd.uscourts.gov/cases/show.php?db=special\&id=306>). In Chapter 2 to Chapter 6, I build the IPD problem based on the wholesale contract in which retailers set both e-readers and e-book prices. In Chapter 7, I evaluate the impact of e-books based on the agency contract, which helps link the counterfactual analysis more closely to the e-book policy change in 2010.

⁵The comScore Web Behavior Database captures the detailed browsing and buying behaviors of 100,000 Internet users across the United States. The panel is based on a random sample from a cross-section of more than two million global Internet users who have given comScore explicit permissions to confidentially capture their Web-wide activities. It is weighted so that the distribution of the demographics matches that of the U.S. Internet user population.

time, book title, format (print book or e-book), price, and quantity information.⁶ It also includes demographics such as household income and age (categorized into groups), family size, zip code, etc. Consumers were resampled every year. Among the online shoppers sampled by comScore, 41% bought at least one book a year. There are 20,637 book buyers and 72,619 book purchases over the five-year sample period. Amazon’s market share was 60% on average and increased over time as shown in Figure 4. Hardcover accounts for only 5% of the transactions. I thus group them into “paperbacks” and use “paperbacks” to refer to all print books.

The second data set contains book genre information that I collected from Amazon using web scrapers. For each book title in the first data set, I collect its genre information and prices for both paperback and e-book formats. There are 122,068 pieces of title-format information. I group Amazon’s subgenres into three genres: “lifestyle,” “casual,” and “practical.” Subgenres within the same genre have similar prices, reading purposes, and consumers’ purchasing patterns. In particular, “lifestyle” books usually contain more pictures. “Casual” books usually serve for entertainment purposes. “Practical” books usually require in-depth reading and note-taking.⁷ These features will affect how consumers perceive e-books as substitutes for paperbacks and thus are relevant when grouping subgenres into genres.

Given that households were resampled yearly and that Kindle prices and qualities changed annually in the data, I choose the time period as a year. For each consumer

⁶Book format information is available only for 2011 and 2012. In the demand-side estimation, I integrate over the format choice when calculating the likelihood function for observations in 2008, 2009 and 2010. For instance, the probability of buying q books equals the sum of the following three probabilities: the probability of buying q paperbacks for a Kindle nonowner, the probability of buying q paperbacks for a Kindle owner, and the probability of buying q e-books for a Kindle owner. See the likelihood function section for the details on constructing the latter three probabilities.

⁷(1) “lifestyle” genre includes “lifestyle & home,” “cooking,” “travel,” “fitness & dieting,” “crafts, hobbies & home,” “arts & photography,” and “children’s book,” etc. (2) “casual” genre includes “fiction,” “science fiction,” “humor,” “nonfiction,” and “biographies & memoirs,” etc. (3) “practical” genre includes “computers & technology,” “business & investing,” “medical books,” and “education & reference,” etc. A typical consumer in the sample buys only one or two books a year. As the number of genres G increases, both the number of zero-consumption choices and consumers’ ex-ante book utility (integrated over G error terms) increase by construction. G is chosen to remain representative of the book heterogeneity while avoiding too many zero choices.

in every period, I aggregate their book purchase records to get their genre-format-retailer level book quantity choices in the demand model. For instance, consumer i bought two “casual” e-books and one “practical” paperback in 2010 from Amazon. I also calculate the average book prices at genre-format-retailer level every year using all book titles in the first data set. The sales-weighted and unweighted prices differ by less than 2%. I use the unweighted ones in the estimation. Of all book sales, 47% were “casual” books, 30% were “lifestyle” books, and 23% were “practical” books. “Casual” books were the cheapest, and “practical” books were the most expensive. Average prices of e-books were lower than paperbacks for all genres. A typical paperback cost \$17.66, and a typical e-book cost \$9.72. For 75.2% of the book titles, e-book prices were lower than their paperback prices.

The third data set contains individual-level Kindle purchase records for 2008 through 2012 from comScore. I observe purchase time, Kindle version, price, quantity, and household demographics. Note that households in the first and the third data sets were resampled every year by comScore. Only a limited number of households were observed for more than one year. It implies that I cannot distinguish between first-time device purchase and upgrading. I take a probabilistic view on the Kindle ownership status in the model. Yearly resampling also implies that I cannot always link individuals’ book transactions in the first data set to their device transactions in the third data set unless they happened during the same year. I thus model consumer behaviors by their observed and unobserved types so that for each type, it is still panel data and both the book-side and device-side transactions are observed.

I supplement the individual-level data with other relevant information. First, I impute aggregate offline book sales from online and offline retailer market shares (Bowker’s Books & Consumers report, 2012). I obtain the offline population size from the fraction of consumers who have bought books online (Nielson Online Shopping

Trend report, 2012)).⁸ Second, I obtain the number of e-books available every year in the Kindle Store from a widely cited blog that takes monthly snapshots of Amazon.⁹ The number of available e-books increased from 126,630 in 2008 to 1,429,500 in 2012. Finally, I impute Kindle cost and book wholesale prices from industry reports.¹⁰ The cost of the most popular Kindle version dropped from \$236 in 2008 to \$89 in 2012.

2.3. Observed Pricing and Consumption Patterns

During the sample periods from 2008 to 2012, Amazon annually launched new Kindle generations and cut the prices for the existing ones, indicating a dynamic pricing strategy on the device (Figure 1). Kindle sales increased over time as shown in Table 2. Unlike Kindle pricing, there was no systematic book price change over time, either overall or by genre-format (Figure 2).¹¹

The data also reveal consumers' book purchase patterns. As for quantity, there is considerable heterogeneity in that 13.8% of the consumers comprised 46.8% of the total book purchases. As for genre, the correlation between consumers' book genre consumption and their observed household characteristics is small, suggesting that the genre choice may be better explained by unobserved heterogeneous genre-specific reading tastes. A noteworthy observation is that consumption across individual book titles is highly dispersed; 92.8% of the book's yearly sales were one, and even the


⁸E-commerce constituted 25.1%, 35.1%, and 43.8% of the U.S. trade book sales from 2010 to 2012, respectively. The rest of the book sales came from offline retailers such as large chain bookstores and independent bookstores. Among book buyers, 44% have purchased books online.

⁹See <http://ilmk.wordpress.com/category/analysis/snapshots/>.

¹⁰The paperback wholesale price is \$15, and the e-book wholesale price is \$12. I impute the Kindle costs from firms that release teardown reports almost every year (<http://www.isuppli.com/Teardowns/News/Pages/Amazon-Kindle-Fire-Costs-201-70-to-Manufacture.aspx>). For years without these reports, I extrapolated data by assuming that the cost drops at the same rate as that of computer parts.


¹¹Online e-book price trackers (e.g., tracker.kindlenationdaily.com) show that there was no systematic price change at the individual title level either. For print books, publishers generally use new book editions to conduct dynamic pricing; they first launch hardcovers that are more expensive and then launch paperbacks and mass market paperbacks. However, they do not dynamically price a particular edition.

Figure 1: Observed Kindle and E-Book Prices



Notes: The graph shows the observed Kindle prices (left y-axis) and observed average e-book prices (right y-axis).

Figure 2: Book Price by Genre and Format



Notes: The graph shows the sales-unweighted prices. The curve for practical paperback prices is scaled down by \$10 to fit into the same graph.

sample bestseller's yearly sales were only 67 or 0.46% of the total yearly sales. The long-tail sales distribution in the sample is comparable to the industrywide distribution. It suggests that a single book title is unlikely to drive the average book prices. As for format, consumers' choices seem to differ by genre. "Casual" books constitute a disproportionately larger share in e-book format (71%) than in paperback format (44%). For e-book buyers, with probability 76.9% they chose e-format for "lifestyle" books, 96.4% for "casual" books, and 61.6% for "practical" books. There is strong substitution between e-books and paperbacks in the same genre; 98.66% of the households bought a particular genre in at most one reading format. In terms of retailer, Amazon's market share increased from 32% in 2008 to 55% in 2012 at the expense of the sales of other online retailers and offline bookstores.

CHAPTER 3

Model Setup

In this chapter, I first illustrate the pricing incentives using a simple two-period model and then describe the full empirical model.

3.1. A Simple Two-Period Model

Consider a two-period model in which a firm sells durable primary hardware at price P and complementary software at price p^E to a unit mass of consumers. The hardware serves as a gateway product to the software and does not bear any stand-alone value.¹² Consumers are heterogeneous in their tastes for the software. The value of a unit of the software v is uniformly distributed on $[0, 1]$. The utility of the hardware comes from the utility generated by subsequent software consumption $u = \lambda(v - p^E) - P$. The coefficient λ is the quantity of the software.¹³ Consumers and the firm share the same discount factor δ and live for two periods. The marginal costs are assumed to be zero. The firm chooses software and hardware prices $\vec{p}_1 = \{p_1^E, P_1\}$ and $\vec{p}_2 = \{p_2^E, P_2\}$ in periods 1 and 2. Consumers have rational expectations about the firm's

¹²I relax this assumption and allow hardware to bear positive values in the full empirical model, represented by quality dummies.

¹³In the full empirical model, the usage intensity λ is endogenized to be a function of the taste parameter v and the book price p^E . I assume that it is a constant here to keep the analytical solution simple while illustrating the same qualitative results.

pricing policies; their beliefs about prices are consistent with the firm's strategy in equilibrium. The marginal consumer in period 1, v_1^* , is indifferent between buying and waiting:

$$\lambda [v_1^* - p_1^E + \delta (v_1^* - p_2^E)] - P_1 = \delta [\lambda (v_1^* - p_2^E) - P_2] \geq 0$$

where $v_1^* = \frac{P_1 - \delta P_2}{\lambda} + p_1^E$. Consumers who buy hardware and software in period 1 are in the range $[v_1^*, 1]$. Similarly, the marginal consumer in period 2, v_2^* , satisfies $\lambda (v_2^* - p_2^E) - P_2 = 0$, where $v_2^* = \frac{P_2}{\lambda} + p_2^E$. Consumers who buy hardware and consume software in period 2 are in the ranges $[v_2^*, v_1^*]$ and $[v_2^*, 1]$, respectively.

This simple setup captures the main features of the traditional single-product IPD as well as new features of IPD with complementary products. In particular, the firm's target is to first extract the most from high-valuation consumers on the hardware and then appeal to low-valuation consumers while earning the most from the software sales. As in the traditional IPD case, the firm faces a shrinking market and lower average willingness-to-pay for the product over time; both the market size and the consumer mix change. A decrease in P_1 reduces the hardware demand in period 2, changes the optimal P_2 , and in turn changes consumer expectations of P_2 as consumers' beliefs are consistent with the optimal strategy. v_1^* summarizes the mass of consumers remaining in the market at the beginning of period 2 and is the relevant state variable for the pricing problem.

Three features are novel in the complementary product setup. First, consumers are self-selected into buying the hardware based on their heterogeneous tastes for the software. Second, the demands of the two products are interrelated. Consumers trade off between the utility from a current hardware purchase and the value of waiting, both of which further depend on the current and future software prices. Third, the firm needs to coordinate the pricing of the two products. p^E affects the profits from

a hardware owner, while p^E and P jointly affect the number of hardware owners. The full model captures all the features of the simple model while allowing for richer heterogeneity and nonlinear demand elasticities.

Using backward induction to solve for period 2 and period 1 prices, I get

$$\begin{aligned}\vec{p}_2(\vec{p}_1) &= \arg \max_{\vec{p}_2} \pi_2 = (1 - v_2^*) \lambda p_2^E + (v_1^* - v_2^*) P_2 \\ \vec{p}_1 &= \arg \max_{\vec{p}_1} \pi_1 + \delta \pi_2 = (1 - v_1^*) (\lambda p_1^E + P_1) + \frac{\lambda \delta}{4} \left(1 + \frac{1}{\delta} \left(p_1^E + \frac{P_1}{\lambda} - 1 \right)^2 \right)\end{aligned}$$


The optimal prices in period 2 are $P_2^* = \frac{\lambda}{2}$ and $p_2^{E*} = 0$. The optimal prices in period 1 satisfy $p_1^{E*} + \frac{P_1^*}{\lambda} = 1 + \delta$. In particular, $P_1^* = \lambda(1 + \delta)$ and $p_1^{E*} = 0$ if $\lambda \geq 1$, and $P_1^* = 0$ and $p_1^{E*} = 1 + \delta$ if $\lambda < 1$.

The optimal strategies with complementary products differ from the harvesting strategy in the traditional single-product IPD case in two ways. First, both harvesting and investing can be optimal. If $\lambda \geq 1$, it is optimal to harvest on the hardware and invest in the software. The opposite is true for $\lambda < 1$. Similarly in the full model, I find that the firm should harvest on Kindles and invest in e-books for the avid readers with high λ and should do the opposite for the general readers with low λ . Second, the firm needs to coordinate p^E and P . The optimal pricing condition $p_1^{E*} + \frac{P_1^*}{\lambda} = 1 + \delta$ indicates that the optimal P increases as p^E decreases within the same period. The results from the full model echo the results from this simple model.

3.2. Consumer Problem

In this section, I outline consumers' discrete Kindle adoption/upgrading decisions and continuous book purchase decisions. Consumers' decision timing is shown in Figure 3. Every period, consumers first make ex-ante dynamic device decisions given the current Kindle price and quality, book price and e-book availability, their beliefs on the future values of these variables, and idiosyncratic device-side shocks. A consumer

Figure 3: Consumer Decision Timing



who does not have a Kindle chooses to buy one or wait for the next period. The benefit of becoming a Kindle owner is that in addition to buying paperbacks, he can buy e-books that are potentially cheaper, more convenient to read, and become more available over time. He needs to trade-off between the gain in discounted book flow utilities from hereon and the one-time payment of the Kindle price. If he chooses to wait, he cannot gain for now, but he can potentially get better Kindle prices and qualities in the future. Kindle owners choose whether to upgrade to the latest Kindle generation or wait. Upgrading is motivated by a higher Kindle quality and does not affect their book purchases. Given their device-adoption statuses and the idiosyncratic book-side shocks, consumers then make decisions in each genre about book purchases (buy or not), format (paperback or e-book), and retailer (Amazon, other online retailers, or offline bookstores, conditional on buying paperbacks). They never drop out of the market.

Consumers' book consumption is modeled at the genre-format-retailer level instead of book title level for two reasons. First, aggregate book sales are more relevant

in the pricing problem compared with single title sales.¹⁴ Second, I do not have title-level aggregate book sales data and cannot estimate title fixed effects to account for price endogeneity issues. The benefit of modeling at the genre level is that the average genre price is not endogenous to the quality of a particular book title in that genre. I estimate genre fixed effects to capture average genre qualities. Section 4.2 provides more discussion on this.

I assume that consumers have persistent heterogeneous genre-specific book tastes. They respond to time-varying book prices, availability, and idiosyncratic shocks by adjusting their usage intensity or book quantities. Consumer taste segments are fixed, while their device ownership distribution evolves endogenously over time. They have perfect foresight on prices, Kindle quality, and book availability.¹⁵ Kindle launching and book availability are taken from the data. Kindle qualities are taken as given and estimated in the model. For years beyond the sample period from 2008 to 2012, I assume that these variables stop evolving and stay at the year 2012 level.¹⁶ I also make the following assumptions for tractability and data limitation reasons. First, I assume that consumers read only e-books on e-readers and not on other screens such as PCs and tablets; that is, they need to buy an e-reader before purchasing and reading e-books. I conduct robustness checks by allowing consumers to buy or read on other devices after 2010. The predictions on dynamic pricing are qualitatively robust.¹⁷ Second, I assume that consumers use only one Kindle at a time and that

¹⁴Meanwhile, modeling at the title-level would require strong assumptions on the books that enter consumers' choice set. It is not appealing to assume that consumers must decide from the millions of books that are available. It is also not appealing to assume that consumers consider only bestsellers, as the majority of the book titles purchased were not bestsellers; 99.94% of the titles were purchased less than 10 times in the data. Book consumption is much more dispersed than other content products such as movies and video games.

¹⁵I assume perfect foresight because Amazon changed prices annually in the five-year period, which leads to a short panel. I also try another rational expectation assumption for which consumer exceptions follow an AR(1) process and the coefficients in the AR(1) model are empirically estimated. The results are robust.

¹⁶As a validation of this assumption, Kindle prices experienced a significant drop from 2007 to 2011 and have remained in the \$139-\$199 range since 2011.

¹⁷In the first robustness check, I allow consumers to buy other reading devices in the demand estimation. The estimated Kindle qualities are smaller, while the key demand-side results remain

there is no resale value for Kindles. I also assume that Amazon offers one Kindle version per period, which is the most popular version of the Kindle each period in the data.¹⁸

Book quantity and format and paperback retailer choices. Consumers choose book quantity and format in three genres by maximizing a quadratic direct utility every period. If the consumer buys paperbacks, he decides whether to buy from Amazon.com, other online retailers, or offline bookstores. Index the three genres, “lifestyle,” “casual,” and “practical,” by $g = 1, 2, 3$, respectively. Let subscript i denote consumer type. Let superscript E and P denote e-books and paperbacks, and let $\{p_{gt}^P, p_{gt}^E\}$ denote their prices, respectively. Let superscript 0 denote Kindle nonowners and let q_{igt}^{P0} denote their paperback quantity choices. Let superscript 1 denote Kindle owners, and let $\{q_{igt}^{P1}, q_{igt}^E\}$ denote their paperback and e-book quantity choices. A Kindle owner of type i maximizes his period utility from both paperbacks and e-books, and a Kindle nonowner maximizes his period utility from only paperbacks:

$$\begin{aligned}
\max_{\{q_{igt}^{P1}, q_{igt}^E\}_g} \quad & u_{it}^{book,1} = z + \sum_g \frac{1}{b_i} \left(a_{igt}^P q_{igt}^{P1} + a_{igt}^E q_{igt}^E - \frac{(q_{igt}^{P1} + q_{igt}^E)^2}{2} \right) \\
s.t. \quad & \sum_g (p_{gt}^P q_{igt}^{P1} + p_{gt}^E q_{igt}^E) + z \leq y_i \\
\max_{\{q_{igt}^{P0}\}_g} \quad & u_{it}^{book,0} = z + \sum_g \frac{1}{b_i} \left(a_{igt}^P q_{igt}^{P0} - \frac{(q_{igt}^{P0})^2}{2} \right) \\
s.t. \quad & \sum_g p_{gt}^P q_{igt}^{P0} + z \leq y_i
\end{aligned} \tag{3.1}$$

the same. In the second robustness check, I account for reading e-books on other devices by adding book profits generated on other devices to Amazon’s profit function. I find that in this case, Amazon has weaker incentives to set low e-book prices to induce Kindle adoption because some consumers already own other devices. Yet the joint IPD strategy is not qualitatively changed.

¹⁸In practice, consumers are offered up to two generations of Kindles every year except for year 2012 when three generations were on the market. The most popular version comprised at least 70% of the sales every year. Also, multiproduct firm pricing is computationally prohibitive. Goettler and Gordon (2011) also make this single-product assumption for computational reasons.

Table 1: Optimal Book Quantity Solutions

Quantity choice	Conditions	Intuitions
Kindle owner: $\{q_{igt}^{P1*}, q_{igt}^{E*}\}_g =$		
$\{0, 0\}$	$p_{gt}^P > \frac{a_{igt}^P}{b_i}, p_{gt}^E > \frac{a_{igt}^E}{b_i}$	Neither is worth buying.
$\{a_{igt}^P - b_i p_{gt}^P, 0\}$	$p_{gt}^P < \frac{a_{igt}^P}{b_i}, p_{gt}^E > \frac{a_{igt}^E}{b_i}$ or $p_{gt}^P < \frac{a_{igt}^P}{b_i}, p_{gt}^E < \frac{a_{igt}^E}{b_i},$ $a_{igt}^P - b_i p_{gt}^P > a_{igt}^E - b_i p_{gt}^E$	Only paperback is worth buying, or both are worth buying, but paperback is more attractive.
$\{0, a_{igt}^E - b_i p_{gt}^E\}$	$p_{gt}^P > \frac{a_{igt}^P}{b_i}, p_{gt}^E < \frac{a_{igt}^E}{b_i}$ or $p_{gt}^P < \frac{a_{igt}^P}{b_i}, p_{gt}^E < \frac{a_{igt}^E}{b_i},$ $a_{igt}^P - b_i p_{gt}^P < a_{igt}^E - b_i p_{gt}^E$	Only e-book is worth buying, or both are worth buying, but e-book is more attractive.
Kindle nonowner: $q_{igt}^{P0*} =$		
$\{a_{igt}^P - b_i p_{gt}^P\}$	$p_{gt}^P < \frac{a_{igt}^P}{b_i}$	Paperback is worth buying.
$\{0\}$	otherwise	Paperback is not worth buying.

where a_{igt}^P and a_{igt}^E are heterogeneous book taste parameters. I parameterize them later. b_i can be interpreted as a heterogeneous price coefficient because it enters the optimal quantity choice linearly in front of the price. z is the numeraire, and y_i is household income. The numeraire price is normalized to be 1. The optimal quantity choices in each genre for owners and nonowners are summarized in Table 1.

The quadratic utility functional form has the following advantages compared with the discrete choice logit utility or the constant elasticity of substitution (CES) utility: 1) it allows for multiple-unit and corner solutions (zero consumption) that are common in book purchase patterns; and 2) the optimal quantity solution given this utility form is a linear function of prices.¹⁹ The utility form I assume here implies

¹⁹More flexible quadratic utility specifications yield qualitatively the same demand-side predictions, although the optimal quantity solutions are more complex. Economides, Seim, and Viard (2008) adopt a similar quadratic functional form without allowing for substitution. For a nice survey on direct utility models of consumer choice in marketing, see Chandukala, Kim, Otter, Rossi, and Allenby (2008).

that utilities from different book genres do not interact and that there is perfect substitution between paperbacks and e-books of the same genre.²⁰ These properties are consistent with observed book consumption patterns.

Once a consumer chooses to buy paperbacks, he decides among buying from Amazon.com, other online retailers, or offline bookstores. Sales of online retailers have been growing steadily as part of the total retail sales even before the introduction of e-books.²¹ I use a discrete choice logit structure to parsimoniously capture this trend while allowing for e-reading to influence it. Denote the three retailers by A , B , and O . The retailer utilities contain genre-specific retailer fixed effects and linear time trends. It is $u_{igt}^{retailer,A} = A_{0g} + A_{1g} \cdot t + A_2 \cdot 1\{owner\} + \zeta_{igt}^A$ for Amazon, $u_{igt}^{retailer,B} = B_{0g} + B_{1g} \cdot t + \zeta_{igt}^B$ for other online retailers, and $u_{igt}^{retailer,O} = \zeta_{igt}^O$ for offline bookstores. Offline bookstores serve as the baseline choice, and its fixed effect and time trend are normalized to zero. $\{\zeta_{igt}^A, \zeta_{igt}^B, \zeta_{igt}^O\}$ are i.i.d. logit errors. In particular, I allow Kindle ownership to affect the probability of buying paperbacks from Amazon, which is captured by A_2 .²² The retailer fixed effects help fit the observed retailer market share. The time trends help fit the variation of the shares over time. A_2 helps capture Amazon’s incremental market share that occurred after the introduction of Kindles and e-books.

Consumer heterogeneity. I parameterize the taste parameters to be func-

²⁰By construction, this model specification cannot generate positive numbers of books bought for both formats in the same genre. In the data set, only 1.34% of consumers buy positive quantities of both formats. For these consumers, I assume that there are two shopping occasions in a period. The observation $\{q_{igt}^{P1}, q_{igt}^{E}\}$ is treated as two independent observations, $\{q_{igt}^{P1}, 0\}$ and $\{0, q_{igt}^{E}\}$. This would mildly overestimate the substitution between paperbacks and e-books.

²¹See http://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf.

²²I model this retailer choice as a separate decision from book quantity-format choice because I observe that consumers who choose different retailers do not exhibit different quantity-format choice patterns. It indicates that the two choices do not seem to interact. Alternatively, one can nest the format-quantity choice inside the retailer choice as $u_{igt}^{retailer} = v_{it}^{book}(p^{retailer}) + \text{retailer FE} + \text{time trend} + \zeta_{igt}$, where the retailer utility contains the indirect book utility as a function of the quantity-format choice and retailer-specific paperback prices. This specification is equivalent to my case because the observed paperback prices are generally the same across retailers. It means that the indirect book utility, v_{it}^{book} , is the same for all retailers and will be canceled out in the retailer choice probability. If the observed retailers’ prices are different or if the counterfactual simulation requires changing the retailer paperback prices, the alternative specification can be of interest.

tions of observed demographics, unobserved book reading tastes, and idiosyncratic taste shocks:

$$\begin{aligned}
a_{igt}^P &= \theta_{ig} + \beta_1 D_i^{age} + \eta_{igt}^P \\
a_{igt}^E &= \underbrace{\theta_{ig} + \beta_1 D_i^{age}}_{\text{Baseline taste}} + \underbrace{(\theta_g^E + \beta_2 D_i^{age} + \beta_3 \log n_t^E) \cdot 1\{ebook\}}_{\text{E-format taste}} + \eta_{igt}^E
\end{aligned} \tag{3.2}$$

First, the two formats in the same genre share the same baseline taste. Consumers are heterogeneous in the genre-specific baseline taste or the genre fixed effect θ_{ig} ; some might enjoy reading “casual” books, while others might enjoy reading “practical” books. I model the baseline taste heterogeneity using a finite mixture specification. Second, consumers have genre-specific e-format taste θ_g^E ; some genres might be more suitable for e-reading than others. The e-format taste can vary over time as the number of e-books available increases.²³ Both the baseline taste and e-format taste can vary by age D_i^{age} . Senior consumers generally read more books, are less tech-savvy, and can have lower e-format tastes. Finally, consumers receive idiosyncratic taste shocks η_{igt} that are assumed to be i.i.d. normally distributed with mean zero and standard deviation σ .²⁴

Note that all the variables in the taste parameter a_{igt} linearly affect book demand because $q_{igt}^* = a_{igt} - b_i p_{gt}$. For instance, consumers who have a higher unobserved taste for “casual” books will buy more “casual” books. The price affects the demand through b_i . I allow the price coefficients to vary across income groups $b_i = b_0 +$

²³The number of e-books available is taken as given from the data. It is not directly correlated with e-reader userbase for two reasons. First, publishers’ e-book launching decisions hinge on their concern about the impact of e-format on print format sales. They decide whether to launch the e-book version on a title-by-title basis and change their launching strategy from time to time. Second, unlike the video game industry, e-book introduction is often not retailer-exclusive. Amazon.com and Barnesandnoble.com both had around three million e-books available by 2012, but their e-reader userbases differed by five times. A caveat is that there is no aggregate shock in the model. The e-book availability coefficient may capture some unobserved trend, if any, other than those retailer time trends I have controlled for.

²⁴In a robustness check, I allow η_{igt}^P and η_{igt}^E to be correlated within the same genre. The implied substitution patterns and price elasticities are very robust with respect to this specification change.

$b_1 D_i^{income}$. The shocks η_{igt} in a_{igt} can capture individual price and taste deviations from the average price p_{gt} and average taste of the type. In all, consumers differ in their observed types $\{D_i^{income}, D_i^{age}\}$, unobserved types θ_{ig} , and idiosyncratic shocks η_{igt} . The unobserved types and the observed types are independent. Consumers have persistent unobserved genre-specific tastes θ_{ig} . They respond to time-varying prices, e-book availability, and idiosyncratic shocks by changing how many books to buy.

As shown in the estimation results later, the unobserved heterogeneous genre fixed effect θ_{ig} is the only difference between avid and general readers. It is the major heterogeneity that drives the difference in consumer book and e-reader consumption and helps fit the genre market share. Paperbacks and e-books differ in prices, consumers' format tastes, and taste shocks. Substitution between the two formats in each genre is explained by the e-format genre fixed effects θ_g^E and the format price difference. Variation in the substitution patterns over time is mainly explained by the time-varying e-book availability n_t^E .

Indirect flow utility from books. I can calculate the ex-ante indirect flow utilities from books for Kindle nonowners and owners, $v_{it}^{book,0}$ and $v_{it}^{book,1}$, by substituting the optimal book quantities into the utility function and taking expectations over the error terms η_{igt} in a_{igt} :

$$\begin{aligned}
v_{it}^{book,0} &= y_i + \sum_g E\left[\frac{(a_{igt}^P - b_i p_{gt}^P)^2}{2b_i} \mid q_{igt}^{P0*} > 0\right] \cdot \Pr(q_{igt}^{P0*} > 0) \\
v_{it}^{book,1} &= y_i + \sum_{F=\{P,E\}} \sum_g \left\{ E\left[\frac{(a_{igt}^F - b_i p_{gt}^F)^2}{2b_i} \mid q_{igt}^{F1*} > 0, q_{igt}^{-F1*} = 0\right] \right. \\
&\quad \left. \cdot \Pr(q_{igt}^{F1*} > 0, q_{igt}^{-F1*} = 0) \right\}
\end{aligned} \tag{3.3}$$

Device adoption decision. Dynamically, given the utilities from books, consumers decide ex ante whether to buy or upgrade their Kindles. Equation 3.4 shows the flow utilities for a Kindle nonowner who chooses to wait, a Kindle owner who

chooses to wait, and a consumer who chooses to buy/upgrade:

$$\begin{aligned}
\bar{u}_{it}^0 &= \Gamma v_{it}^{book,0} + \bar{\varepsilon}_{it}^0 \\
\bar{u}_{it}^1 &= \Gamma v_{it}^{book,1}(p_t^E) + \bar{Q}_{it} + \bar{\varepsilon}_{it}^1 \\
u_{it} &= \Gamma v_{it}^{book,1}(p_t^E) + Q_t - \alpha_1 P_t + \varepsilon_{it}
\end{aligned} \tag{3.4}$$

If a consumer does not have a Kindle and chooses to wait in this period, he receives book utility only from paperbacks. If he has a Kindle and chooses not to upgrade, he receives book utility from both paperbacks and e-books and the quality of his old Kindle $\bar{Q}_{it}$. If a consumer chooses to buy/upgrade to the latest Kindle generation, he receives book utility from both paperbacks and e-books plus the new Kindle quality Q_t at the cost of Kindle price P_t . The price coefficient $\alpha_i = \alpha_0 + \alpha_1 D_i^{income}$ is allowed to vary across income groups. The idiosyncratic shocks $\{\bar{\varepsilon}_{it}^0, \bar{\varepsilon}_{it}^1, \varepsilon_{it}\}$ are identically and independently distributed extreme value type I errors, which are also independent of the error terms on the book side. The variances are normalized to be 1.²⁵ The Kindle qualities are estimated as dummies. I do not allow them to interact with book utilities because the data cannot identify such interaction. When adding Kindle quality dummies into the book utility, the estimated dummies are insignificant. The reason might be that, for instance, a Kindle owner who had a new Kindle in 2012 did not buy significantly different numbers of e-books compared with a Kindle owner who had an old generation in 2008. I thus keep the Kindle qualities and book utilities as additively separable in the device flow utility.

The flow utility enters the dynamic programming problem. To make the notation more general, I use $\bar{u}_{it}$ to jointly denote the flow utility of waiting for nonowner and owner $\{\bar{u}_{it}^0, \bar{u}_{it}^1\}$. For a nonowner, Kindle quality is $\bar{Q}_{it} = 0$. For each consumer type,

²⁵Another normalization approach is to drop the coefficient Γ and estimate the variance of the error term. The two approaches are equivalent. The coefficient Γ is thus identified by the observed variation in the Kindle adoption decisions, similar to how variance of the error term is identified in the second normalization approach.

the state space contains (1) the current Kindle ownership status $\bar{Q}_{it}$, which evolves based on the device adoption choice; (2) the e-book price p_t^E , which enters the ex-ante flow utility from books v_{it}^{book} ; (3) the offered Kindle price P_t and quality Q_t ; and (4) the idiosyncratic shocks on the device side $\vec{\varepsilon}_{it} \equiv \{\bar{\varepsilon}_{it}, \varepsilon_{it}\}$. Let $V(\bar{Q}_{it}, Q_t, P_t, p_t^E, \vec{\varepsilon}_{it})$ denote the value function of a consumer with current device $\bar{Q}_{it}$ at the beginning of the period. $d_{it} = 1$ indicates buying/upgrading and $d_{it} = 0$ indicates waiting. The Bellman equation is

$$V(\bar{Q}_{it}, Q_t, P_t, p_t^E, \vec{\varepsilon}_{it}) = \max \left\{ \bar{u}_{it} + \delta E[V(\bar{Q}_{it}, Q_{t+1}, P_{t+1}, p_{t+1}^E, \vec{\varepsilon}_{it+1}) \mid Q_t, P_t, p_t^E, d_{it} = 0], \right. \\ \left. u_{it} + \delta E[V(Q_t, Q_{t+1}, P_{t+1}, p_{t+1}^E, \vec{\varepsilon}_{it+1}) \mid Q_t, P_t, p_t^E, d_{it} = 1] \right\} \quad (3.5)$$

The first and second elements of the max operator are the choice-specific value functions of waiting and buying/upgrading. Conditional on waiting, the device adoption status remains at $\bar{Q}_{it+1} = \bar{Q}_{it}$. Conditional on buying/upgrading, the device adoption status evolves deterministically as $\bar{Q}_{it+1} = Q_t$. The rest of the state space $\{Q_t, P_t, p_t^E\}$ evolves to $\{Q_{t+1}, P_{t+1}, p_{t+1}^E\}$, according to consumers' expectation about next period values $h(Q_{t+1}, P_{t+1}, p_{t+1}^E \mid Q_t, P_t, p_t^E)$. Let $EV(\cdot) = \int_{\vec{\varepsilon}} V(\cdot, \vec{\varepsilon}) dg_{\vec{\varepsilon}}$ denote the expectation of the value function integrated over $\vec{\varepsilon}_{it}$. The expected value function is

$$EV(\bar{Q}_{it}, Q_t, P_t, p_t^E) = \ln \left[\exp(\bar{u}_{it} - \bar{\varepsilon}_{it} + \delta E[V(\bar{Q}_{it}, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 0]) \right. \\ \left. + \exp(u_{it} - \varepsilon_{it} + \delta E[V(Q_t, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 1]) \right] \quad (3.6)$$

Notice that there is a unique expected value function for each type. The probability

vector of buying/upgrading for each type is

$$\begin{aligned}\phi(d_{it} = 1 \mid \bar{Q}_{it}, Q_t, P_t, p_t^E) &= \frac{A}{A + B} \\ A &= \exp(\bar{u}_{it} - \bar{\varepsilon}_{it} + \delta E[V(\bar{Q}_{it}, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 0]) \\ B &= \exp(u_{it} - \varepsilon_{it} + \delta E[V(Q_t, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 1])\end{aligned}$$

where each element in the vector represents the probability of adoption/upgrading for consumer type i with current Kindle quality $\bar{Q}_{it}$.

The key feature of the demand system is that Kindle adoption is driven by usage intensity of books, which is further endogenized to be a function of consumers' book tastes and book prices. In this sense, e-book prices affect Kindle attractiveness. The book-side and device-side decisions are linked because (1) the ex-ante flow utilities from books affect the Kindle adoption decisions; and (2) the Kindle adoption statuses influence the book formats that consumers can choose from. Consumers are motivated to buy Kindles for three reasons: the gain from current-period book utility, the current device prices and qualities, and the option value of device adoption. To see this, take the difference of the two choice-specific value functions:

$$\begin{aligned}&\{[\Gamma v_{it}^{book} + Q_t] - [\Gamma v_{it}^{book} + \bar{Q}_{it}]\} - \alpha P_t + \\&\delta\{E[V(Q_t, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 1] \\&- E[V(\bar{Q}_{it}, Q_{t+1}, P_{t+1}, p_{t+1}^E) \mid Q_t, P_t, p_t^E, d_{it} = 0]\}\end{aligned}$$

The first term represents the increase in Kindle quality and, for a first-time device adopter, the increase in book flow utility (v_{it}^{book} changes from $v_{it}^{book,0}$ to $v_{it}^{book,1}$). The second term indicates that consumers will respond to a Kindle price drop. The third term is the option value, or the discounted utility gain from Kindles in the future. Both the current and future gains drive consumers to self-select into buying Kindles.

We can expect that consumers who like reading benefit more from having Kindles and will adopt earlier.

3.3. Firm Problem

I take a normative stance and estimate the demand system without assuming optimality of the observed prices. Given Kindle costs and book wholesale prices from industry reports, I use the estimated demand system to compute Amazon’s optimal Kindle and e-book pricing strategies. The consumers’ and the firm’s dynamic problems are jointly solved in the simulation. The book wholesale prices are adjusted by a calibrated spillover effect to Amazon’s other product business, so as to reflect a more realistic marginal cost of selling books for Amazon.²⁶ To keep the model tractable, I make several simplifications from the demand model. First, I abstract from the quality improvements to analyze IPD because I do not have data on R&D. It is also computationally prohibitive to jointly solve for quality and pricing problems with complementary products. Because quality improvements are often intertwined with price changes, I use the average estimated quality in the pricing simulation. Consumers have only two device ownership statuses: Kindle owner and nonowner.²⁷

²⁶I first impute the book wholesale prices from the standard pricing approach in the publishing industry. The list price of e-books is 80% of the list price of paperbacks. Amazon sells books at 60% of the list price on average. The wholesale price for both paperbacks and e-books is 50% of the list price. I use these rules and the observed Amazon paperback prices to back out the wholesale prices, which are \$15 for paperbacks and \$12 for e-books. Yet the observed Amazon e-book price is \$9.72, and the simulation cannot generate such a low price level. The reasons might be that there are unobserved factors that change Amazon’s actual marginal cost such as spillover effect into Amazon’s other product business, negotiated quantity discounts that are not publicly observed, and competition pressure. To get a more realistic marginal cost value, I allow for a spillover effect per book transaction in the simulation. I solve the pricing problem with different magnitudes of this spillover effect. The predictions on dynamic pricing are very robust, which can be regarded as robustness checks on the value of the marginal cost I choose. I pick one value of the spillover effect so that the simulated e-book price *level* is comparable to the observed one for just the first period. Notice that this does not match the entire price path because I still take a normative view on the dynamic pricing policy. I keep this spillover effect when reporting the pricing and profitability results. Gentzkow (2007) adopts a similar approach when rationalizing the zero price of online newspapers.

²⁷The upgraders are modeled in a simplified way. They have proportionally higher book flow utilities than first-time adopters. The proportion is calculated from the estimated demand system.

Second, I restrict the pricing policy to be functions of only the two unobserved types and average over the observed demographics types. This helps reduce the state space greatly from 36 dimensions to two dimensions while keeping the major heterogeneity in consumer tastes that drives self-selection. Third, I do not distinguish between book genres and solve for one e-book price as if changing book prices uniformly across genres.

Amazon sets Kindle price P and e-book price p^E to maximize its total discounted profits from Kindles, print books, and e-books. The firm's state space Δ is a vector that contains the number of Kindle nonowners for each type at the beginning of the current period, similar to v_1^* in the simple two-period model. The demand system provides two key inputs to the firm's pricing problem: (a) the probabilities of adopting Kindles $\phi_t = \{\phi_{it}\}_i$; (b) the book profits Amazon earns from each Kindle owner $r_{it}^1 = (p_t^E - w^E) \cdot q_{it}^E(p_t^E) + (p^P - w^P) \cdot q_{it}^{P1A}(p_t^E)$ and nonowner $r_{it}^0 = (p^P - w^P) \cdot q_{it}^{P0A}$, where $\{w^P, w^E\}$ are the wholesale prices paid to publishers and $\{q_{it}^{P1A}, q_{it}^{P0A}\}$ are the numbers of paperbacks sold on Amazon, which equal the total number of paperbacks $\{q_{it}^{P1}, q_{it}^{P0}\}$ times the probability of buying from Amazon $\Pr(\text{Amazon} \mid q_{it}^P > 0, q_{it}^E = 0)$. Let $r_t^1 = \{r_{it}^1\}_i$ and $r_t^0 = \{r_{it}^0\}_i$. Let $r_t^{gain} \equiv r_t^1 - r_t^0$ denote Amazon's gain in book profits from converting a nonowner to an owner. Given the consumer adoption probability ϕ_t , the number of nonowners in the next period Δ_{t+1} equals the probability of not buying/upgrading times the number of nonowners in this period Δ_t , which indicates that the state space evolves deterministically as $\Delta_{t+1} = [I - \phi_t] \Delta_t$. The Bellman equation of the firm is

$$\begin{aligned} EW_t(\Delta_t) &= \max_{P_t, p_t^E} \pi_t(P_t, p_t^E, \Delta_t) + \delta E[W_{t+1}(\Delta_{t+1}) \mid P_t, p_t^E, \Delta_t] \\ \pi_t(P_t, p_t^E, \Delta_t) &= \pi_t^{Kindle} + \pi_t^{book,1} + \pi_t^{book,0} \end{aligned}$$

where Kindle profits equal the number of new adopters times the Kindle mark-up

$\pi_t^{Kindle} = (\phi_t \cdot \Delta_t) [P_t - c_t]$. Aggregate book profits from Kindle nonowners equal the cumulative number of nonowners times the book profits generated per nonowner $\pi_t^{book,0} = (I - \phi_t) \cdot \Delta_t \cdot r_t^0$. Aggregate book profits from Kindle owners equal the cumulative number of owners times the book profits generated per owner $\pi_t^{book,1} = (\Delta_0 - \Delta_t + \phi_t \cdot \Delta_t) \cdot r_t^1$, where Δ_0 is the initial market size at period 0. Note that I need to compute the value function for each period because the following variables change over time: (1) Kindle cost c_t , which affects π_t^{Kindle} and is taken from the data; and (2) consumer's retailer choice probability $\Pr(\text{Amazon} \mid q_{it}^P > 0, q_{it}^E = 0)$, which affects $\{r_t^0, r_t^1\}$ and changes over time based on the time trends in the retailer utility. Same as in the demand model, I assume that these variables stop evolving and remain at the year 2012 level to keep the problem stationary.

Taking the F.O.C. with respect to the Kindle price yields

$$\underbrace{\Delta_t \cdot \frac{\partial \phi_t}{\partial P_t} [P_t - c_t] + \phi_t \cdot \Delta_t}_{\text{static Kindle profit change}} + \underbrace{\Delta_t \cdot \frac{\partial \phi_t}{\partial P_t} \cdot r_t^{gain}}_{\text{static book profit change}} - \underbrace{\Delta_t \cdot \frac{\partial \phi_t}{\partial P_t} \cdot \delta \frac{\partial W_{t+1}(\Delta_{t+1})}{\partial \Delta_{t+1}}}_{\text{dynamic future state change}} = (37)$$

The first-order condition informs the trade-offs for Kindle pricing. Statically, the firm needs to get a greater number of Kindle owners (affected by $\{p_t^E, P_t\}$) and earn higher profits from each owner (affected by p_t^E). A higher Kindle price increases the marginal gain on the existing Kindle sales (the first term) at the cost of the gains from new adopters (the second term) and their associated book profits (the third term). The demand elasticities dictate the magnitudes of these effects. Dynamically, two effects are captured in the fourth term: (1) a higher current Kindle price reduces the market size and changes the future mix of the two consumer types; and (2) the current prices today affect consumers' expectation over future prices, which in turn affects current adoption. The firm needs to tradeoff between static and dynamic incentives and optimally manage the size and mix of owners as well as profits per owner. Taking the F.O.C. with respect to the e-book prices yields the following equation with similar

trade-offs:

$$\begin{aligned}
& \underbrace{\Delta_t \cdot \frac{\partial \phi_t}{\partial p_t^E} \cdot r_t^{gain} + [\Delta_0 - \Delta_t + \phi_t \cdot \Delta_t] \cdot \frac{\partial r_t^{gain}}{\partial p_t^E}}_{\text{static book profit change}} + \\
& \underbrace{\Delta_t \cdot \frac{\partial \phi_t}{\partial p_t^E} [P_t - c_t]}_{\text{static Kindle profit change}} - \underbrace{\Delta \cdot \frac{\partial \phi_t}{\partial p_t^E} \cdot \delta \frac{\partial W_{t+1}(\Delta_{t+1})}{\partial \Delta_{t+1}}}_{\text{dynamic future state change}} = 0
\end{aligned} \tag{3.8}$$

I consider the pure-strategy Markov-perfect Nash equilibrium (MPNE) in which both consumers and the firm are forward-looking. The noncommitment pricing policy is subgame perfect in that prices are optimal given the state of the market in any period.²⁸ The setup is similar to the frameworks in Nair (2007) and Goettler and Gordon (2011). The equilibrium requires that the consumer's expectation over the future state is consistent with the firm's optimal strategy.²⁹ The equilibrium is defined as the set $\{V^*, W^*, P^*, p^{E*}, h^*\}$, which contains the equilibrium value functions for the consumers and the firm, the optimal pricing policy functions for Kindles and e-books, and the beliefs about next period state space.³⁰

²⁸I consider noncommitment policies because “policies with commitment are not generally subgame perfect” and “the firm has an incentive to deviate from the announced policy after the initial period passes” (Besanko and Winston, 1990). Noncommitment policies are “more managerially relevant” (Nair, 2007).

²⁹Consumer rational expectation assumption is common in dynamic equilibrium models. On the theoretical side, such “relatively simple equilibrium policies are effective in explaining the key qualitative features of the data” (Nair, 2007). In practice, the prices of digital durable goods such as iPhones drop in a regular manner. Many online websites (e.g., decide.com) also provide consumers with price drop predictions based on historical prices, which further enhance consumers' ability to predict price change.

³⁰When solving for the equilibrium, I need to jointly solve the consumers' and the firm's dynamic problems. In equilibrium, the Kindle price $P = P(\Delta)$ and the e-book price $p^E = p^E(\Delta)$ are functions of the state space. In the simulation, it is useful to rewrite the consumer's problem with Δ in the state space instead of P and p^E . The simplifications on the supply side (e.g., book prices change uniformly across genres) apply on the demand side as well.

CHAPTER 4

Demand Estimation and Supply Simulation Method

This chapter constructs the likelihood function based on the demand model and discusses identification, demand estimation, and supply simulation strategies.

4.1. Likelihood Function

The total log likelihood is composed of probabilities of the individual-level device and book choice observations, as well as aggregate offline book sales $L = L^{Kindle} + L^{book} + L^{aggregate}$.

On the device side, the device choice gives the conditional probability of buying/upgrading conditional on holding Kindle version $\bar{Q}_{it}$. It is $\phi(d_{it} = 1 \mid \bar{Q}_{it}, Q_t, P_t, p_t^E)$ or $\phi(d_i = 1 \mid \bar{Q}_{it}, t)$, as $\{Q, P, p^E\}$ are unique per period. The conditional probability of buying/upgrading further implies the probabilities of holding a particular Kindle generation at time t , $\Pr(\bar{Q}_{it} \mid t)$. Combining these two probabilities, I can calculate the unconditional probabilities of buying/upgrading every period $\Pr(d_i = 1 \mid t) = \sum_{\bar{Q}_{it}} \phi(d_i = 1 \mid \bar{Q}_{it}, t) \Pr(\bar{Q}_{it} \mid t)$. The device part of the log likeli-

hood function for each consumer type is

$$L_i^{Kindle} = \sum_{t=2008}^{2012} [n_{i1t} \log [\Pr (d_i = 1 | t)] + n_{i0t} \log [1 - \Pr (d_i = 1 | t)]]$$

Here n_{i1t} is the observed Kindle sales from type i at time t , and $n_{i0t} = N_{i0} - n_{i1t}$ is the observed number of waiting decisions. N_{i0} is the number of type i consumers in the initial market. Summing over the observed types and integrating over the unobserved types, I can obtain the total device-side log likelihood function L^{Kindle} .

On the book side, the individual-level probabilities of book quantity, format, and paperback retailer choices are combined to form the likelihood L^{book} . For instance, a Kindle owner i buys q paperbacks from Amazon and zero e-book in genre g at time t . The probability of this observation is

$$\underbrace{f(q_{igt}^P = q > 0, q_{igt}^E = 0)}_{\text{format=P, quantity=q}} \underbrace{\Pr(\text{Amazon} | q_{igt}^P > 0, q_{igt}^E = 0)}_{\text{retailer=Amazon}}$$

Taking log and summing over all the observations for each individual every period, I can obtain the total book-side log likelihood function L^{book} .

Finally, I match the model-predicted aggregate offline book sales $\hat{H}_{gt}(\Omega)$ to the observed ones in the data H_{gt} . Given parameter values Ω , I simulate the error terms in the taste parameter $\{\eta_{igt}^P, \eta_{igt}^E\}$ 10,000 times and calculate the predicted offline book quantity for individual i in genre g at time t . Summing over the individual quantities, I obtain the predicted aggregate offline book sales $\hat{H}_{gt}(\Omega)$. The set of simulated error terms is fixed throughout the estimation to keep the problem stationary. Following Allen, Clark, and Houde (2014), I calculate the probability of observing $\hat{H}_{gt}(\Omega)$ using the central-limit theorem. The null hypothesis is that the model is correctly specified so that $\hat{H}_{gt}(\Omega) - H_{gt}$ is normally distributed with mean zero and variance σ_{gt}^2 , where σ_{gt}^2 is the predicted variance of $\hat{H}_{gt}(\Omega)$. The likelihood of the aggregate moments

is thus $L^{aggregate} = \sum_{g,t} \log \frac{1}{\sigma_{gt}} \phi \left[\left(\hat{H}_{gt}(\Omega) - H_{gt} \right) / \sigma_{gt} \right]$, where $\phi(\cdot)$ is the PDF of standard normal distribution. The device-side, book-side, and aggregate moment likelihood jointly form the total log likelihood $L = L^{Kindle} + L^{book} + L^{aggregate}$.

4.2. Identification

The book-side parameters include taste parameters $\{\theta_{ig}, \theta_g^E, \beta_1, \beta_2, \beta_3, \sigma\}$, retailer fixed effects and time trends, and price coefficient b_i . I can group them as time-invariant (fixed effects) and time-variant (prices, e-book availability, and time trends). The genre fixed effects in the baseline taste θ_{ig} are captured by a finite mixture model and are identified from the genre market shares. The genre fixed effects in the e-format taste θ_g^E are identified from the substitution patterns between paperbacks and e-books. Variation in the substitution patterns over time identifies the coefficient on time-varying e-book availability β_3 . The coefficients on age β_1 and β_2 are identified from the consumption patterns across age groups. The retailer fixed effects and time trends are identified from retailer market share and its variation over time. The rest of the variation is explained by price. In particular, the price coefficient b_i is identified from the variation in book prices across genres, formats, and over time. The price variation over time is small and comes from exogenous changes in, for instance, the mix of book titles and the institutional pricing arrangements. A larger price variation comes from the difference between paperback and e-book prices across genres. In the model, I control that paperbacks and e-books in the same genre share the same persistent genre baseline taste. This helps identify the price coefficient.³¹

³¹If there were only two genres and prices are constant over time, the genre fixed effects in the baseline taste θ_{ig} and the genre fixed effects in the e-format taste θ_g^E would capture all the difference among the four genre-format combinations, leaving no variation to identify the price coefficient. In my model, there are three genres and the two formats in the same genre share the same baseline taste θ_{ig} . Besides, $p_{gt}^P - p_{gt}^E$ differs by genre. This imposes extra restrictions on the parameters so that the price coefficient is identified (even without price variation over time). The change in $p_{gt}^P - p_{gt}^E$ over time provides further identification sources.

The device-side parameters include $\{\Gamma, \alpha_i, \{Q_t\}_{t=2008}^{t=2012}\}$. The coefficient on book utility Γ is cross-sectionally identified from the different device adoption probabilities of consumers given the same Kindle price and quality. The price coefficient α_i and Kindle quality dummies $\{Q_t\}_{t=2008}^{t=2012}$ are jointly identified from two sources: (1) cross-sectionally, the different adoption/upgrading probabilities of consumers owning different Kindle versions; and (2) intertemporally, the adoption/upgrading probabilities for each consumer type. They are separately identified, as price is incurred only once and quality enters utility every period. I assume that the Kindle quality evolves according to a Markov process $\{x_n\}$ with a degenerate transition matrix. The transition probability is $p(x_{n+1} = Q_{t+1} | x_n = Q_t) = 1$ for $t = 2008$ to 2011 and $p(x_{n+1} = Q_{2012} | x_n = Q_{2012}) = 1$. To avoid overfitting, I capture the values of the qualities using $Q_t = Q_0 + Q_1 \log t$.

Since the book side and the device side are linked by consumers' self-selection process, many parameters are identified from both sides. For instance, the taste parameters and price coefficient on the book side enter book utility and in turn device utility, so they are also identified from device adoption choices. The income-specific price coefficients and age-specific taste coefficients are identified by both the observed book purchases and the observed Kindle adoption patterns across income-age groups.

Upgrading can be identified by observing how e-book sales change as Kindle sales change. In a world without upgrading, additional Kindle sales (times the number of e-books bought per Kindle owner) directly yield additional e-book sales. In a world with upgrading, additional Kindle sales that come from upgrading do not lead to additional Kindle owners and thus do not yield additional e-book sales. Furthermore, the demographics of Kindle buyers over time offer clues about returning consumers. Without upgrading, the change in the demographic composition should be monotonic as the consumer pool is exhausted. Upgraders can be identified if the income and age distribution of the Kindle adopters in the later years are similar to that in the early

years.

Price endogeneity. The demand estimation is conducted without imposing pricing optimality conditions. The observed prices might be endogenized to unobserved qualities and demand shocks. The model setup helps eliminate the price endogeneity issue. First, to account for the fact that Kindle prices might be endogenous to their qualities, I explicitly model and estimate the qualities of different Kindle generations. Second, I model book consumption at the genre level instead of at the book-title level. The price of a particular book title might be endogenous to its quality, which in turn affects its sales. In the model, I use the average genre price, which is not endogenous to the qualities of individual book titles in that genre. In other words, a single book title does not drive the average genre price.³² The genre qualities are estimated as fixed effects in the model. Finally, prices are not endogenous to demand fluctuations over time.

To ensure that best sellers do not drive average genre prices in the data, I tabulate the sales of the book titles per year. It turns out that 92.82% of the book titles had only one purchase record, 5.53% had two purchases, and 99.94% had less than 10 purchases. The consumption is highly dispersed. This pattern holds for all genres and reading formats. The title with the highest sales was bought 67 times among 14,524 total book transactions in that year. This is comparable to the ratio of the best sellers' sales to the total trade volume in the industry. The sales-weighted and unweighted prices on average differ only by 2%, and one of them is not systematically higher or lower than the other. I conduct robustness checks by estimating the model using both sales-weighted and unweighted prices.³³ The results are quite robust.

³²A single book title is not likely to drive consumers' device choice, either. Unlike in the video game and console market, consumers can always buy a title in paperback format, while they cannot play games without buying the console. Furthermore, e-books are mostly not exclusive to retail platforms. Amazon and Barnesandnoble.com have a comparable number of e-books available, while many video games are exclusive to particular consoles.

³³Thanks to Song Yao for suggesting this robustness check.

4.3. Estimation and Computational Methods

To estimate the demand model, I use the Nested Fixed Point algorithm (NFXP) proposed by Rust (1987). For each iteration, I solve the dynamic programming problem for each consumer type in the inner loop and use MLE in the outer loop. Given a set of parameter guesses, I first calculate the book-side probabilities and the flow utilities from books. I then feed the flow utilities to the device side and solve the dynamic programming problem using the value function iteration method. The flow utilities and expected value functions are calculated separately for the 36 consumer types (three age groups times three income groups times four unobserved segments on book tastes). In particular, the value functions need to be calculated for consumers holding different generations of Kindles. Given the value functions, I can construct the device-side probabilities and form the total likelihood function.

To solve for the supply-side pricing problem, note that the equilibrium requires that (1) consumers and the firm make optimal decisions, and that (2) consumer's expectation over the future state is consistent with the firm's optimal strategy. The computation algorithm includes an inner loop and an outer loop, which I detail in the appendix. In the inner loop, Condition (1) requires that consumers solve their maximization problems given their beliefs about next period state space and firm's policy functions. The firm then uses the consumers' choices to update next period state space and solve their policy functions again. Condition (2) requires that I repeat this process until the optimal policy functions and the next period states reach a fixed point. Given the fixed point, the outer loop updates the value function guesses and iterates until convergence.

Function approximations are used in the demand estimation and supply simulation. First, calculating the device choice probability requires calculating the indirect utilities from books $v_{it}^{book,1}$ and $v_{it}^{book,0}$, which contain conditional expectations of a

truncated normal error and its quadratic term. For $v_{it}^{book,1}$, the truncation point is a result of a max operator. I use Gauss-Hermite quadrature with 10 nodes to calculate the conditional expectations. The format-quantity choice probability also involves the probability of a truncated normal error in which the truncation point is a result of a maximization operator. There is no closed-form expression for it. I use Gauss-Chebyshev quadrature with 10 nodes and Gauss-Laguerre quadrature with 10 nodes to approximate the integrals. The details are presented in the appendix. Second, several demand-side variables enter the supply-side problem: the probability of adopting Kindles ϕ and the book profits generated per Kindle owner r^1 and nonowner r^0 . They are functions of Kindle price P and e-book price p^E . I evaluate them on a set of grid points for P and p^E given the estimated demand system. I then approximate these variables as functions of P and p^E using splines.³⁴ Third, I discretize the state space into 20 grid points along each dimension. The value functions are approximated using cubic splines by interpolating between the grid points so that the functions are differentiable when computing the firm's first-order conditions.

³⁴I try both linear splines and cubic splines. Linear splines with 11 breakpoints provide better approximation. This is because the functions are highly linear with little local curvatures and level off as P and p^E increase. Cubic splines produce small fluctuations around the steady value, which make the derivatives inaccurate. The accuracy of the derivatives is important when solving for the firm's first-order conditions. Therefore, I use linear splines to approximate the functions.

CHAPTER 5

Demand-Side Estimation Results

This chapter presents the model fit and discusses the demand estimates and their implications.

5.1. Model Fit

I first present the model fit of both Kindle and book sales at the aggregate level. Table 2 displays the observed and predicted Kindle cumulative sales as a percentage of total market size over time. Figure 4 compares the observed and predicted book sales for both paperbacks and e-books by book genre and by retailer over time. The demand model is estimated using data from 2008 to 2012. Data in 2013 are used as an out-of-sample fit test. The model fits the aggregate-level Kindle and book sales and the trend for each book format, genre, and retailer well.

At the individual level, model predictions can also be validated using survey data. The Pew Research Center conducted surveys in February 2012 and reported that e-reader nonowners bought seven books and owners bought 12 books in the past 12 months. According to the model prediction in Table 4, avid readers' book consumption increases from 12.16 books to 18.26 books once they become a Kindle owner, while general readers' book consumption increases from 1.01 books to 4.96 books. The model predicts that 40% of Kindle owners were avid readers and that 60%

Table 2: Model Fit: Cumulative Kindle Sales/Total Market Size (%)

Year	2008	2009	2010	2011	2012	2013
Observed	0.41	1.36	3.54	6.84	10.34	14.54
Predicted	0.30	1.40	4.02	7.51	11.41	15.73
Standard Error	(0.15)	(0.22)	(0.32)	(0.46)	(0.90)	(1.29)

were general readers by 2012. Overall, the model predicts that a typical Kindle owner buys 7.5 books before and 12.7 books after adopting the device.³⁵ These numbers are comparable to the survey results. Both the aggregate-level and individual-level model fits indicate that the model is able to recover the values of different book formats and retailers, as well as the values of device adopting and waiting.

5.2. Parameter Estimates


Parameter interpretations. Table 3 reports the parameter estimates. Note that the values of the estimated genre fixed effects θ_{ig} and e-format fixed effects θ_g can be interpreted as number of books because they linearly enter consumers’ book taste parameters as $a_{igt} = \theta_{ig} + \beta_1 D_i^{age} + (\theta_g^E + \beta_2 D_i^{age} + \beta_3 \log n_t^E) \cdot 1\{ebook\} + \eta_{igt}$ and optimal quantity choice as $q_{igt}^* = a_{igt} - b_i p_{gt}$.

The estimates on the baseline tastes show that consumers are highly heterogeneous in their unobserved genre-specific reading tastes, which is captured by a finite mixture structure.³⁶ The data reveal four segments. Segments 1, 2, and 3 represent consumers who have high reading tastes for “casual” books, for “lifestyle” and “practical” books, and for all books respectively. They constitute 3.88% of the total

³⁵According to the model estimation, (1) 7.5% of the population are avid readers and 92.5% are general readers, and (2) 38% of avid readers and 2.2% of general readers own a Kindle by 2012. Combining the above estimates, the percentage of Kindle owners who are avid readers equals $(7.5\% \cdot 38\%) / (7.5\% \cdot 38\% + 92.5\% \cdot 2.2\%) = 58\%$ and $100\% - 58\% = 42\%$ for general readers. This means that a typical Kindle owner, based on the model prediction, buys $12.16 \cdot 58\% + 1.01 \cdot 42\% = 7.5$ books before buying a Kindle and $18.26 \cdot 58\% + 4.96 \cdot 42\% = 12.7$ books after adopting one.

³⁶I determine the number of segments by incrementally adding segments until one of the segment sizes is not statistically different from zero. Besanko, Dubé, and Gupta (2003) and Nair (2007) have taken a similar approach. For each genre, I am able to identify two taste levels: high and low. A complete combination of three genres and two levels leads to six types. The estimated segment sizes are significantly different from zero for four out of the six types.

Figure 4: Model Fit: Books



Notes: Observed values are indicated with solid lines, and predicted values are indicated with dashed lines. 95% confidence intervals are shaded. For Graphs 2-4, the lines from top to bottom are sales for Amazon paperbacks, other online retailer paperbacks, and Amazon e-books.

Table 3: Parameter Estimates

Book	Lifestyle: $g = 1$	Casual: $g = 2$	Practical: $g = 3$
Baseline FE			
θ_{ig} : High	10.66*** (0.0025)	11.81*** (0.0028)	12.25*** (0.0188)
θ_{ig} : Low	1.858*** (0.0066)	1.769*** (0.0139)	4.414*** (0.0059)
E-format FE			
θ_g^E	0.9220*** (0.1070)	2.755*** (0.0494)	-1.210*** (0.0134)
Retailer FE & time trends			
A_{0g}	-0.4162*** (0.0091)	-0.7696*** (0.0117)	-0.4627*** (0.0171)
A_{1g} (time)	0.1491*** (0.0044)	0.1380*** (0.0048)	0.2615*** (0.0015)
B_{0g}	-1.529*** (0.0065)	-0.9441*** (0.0101)	-1.6203*** (0.0108)
B_{1g} (time)	0.0198*** (0.0002)	9.317e-4*** (2.003e-5)	0.1368*** (0.0133)

Consumer segment tastes and sizes

Segment	Genre baseline FE $\{\theta_{ig}\}_{g=1,2,3}$	Population mass $\{m_i\}_i$
1	low, high, low	0.0244*** (0.0021)
2	high, low, high	0.0141*** (0.0015)
3	high, high, high	0.0003 (0.0024)
4	low, low, low	$1 - m_1 - m_2 - m_3$

Device	Book				
α_0	0.0048*** (0.0002)	b_0	0.2300*** (0.0026)	β_1	0.1027*** (0.0020)
α_1	-1.003e-4*** (5.040e-6)	b_1	-0.0074*** (0.0017)	β_2	-3.925e-4*** (1.523e-5)
Γ	13.77*** (0.3467)	σ	2.119*** (0.0054)	β_3	0.0050*** (0.0002)
Q_0	-1.068*** (0.0647)	A_2	0.6515*** (0.0267)		
Q_1	0.3429*** (0.0458)				
		MLE Obj.	145,127	# Obs	89,382

Notes: ***, **, * represent significant at the 1, 5, and 10 percent level, respectively.

population, or 7.47% of the book buyers. Segment 4 represents consumers who have low tastes for all genres. The first three segments are much closer along key dimensions of interest such as demand elasticities and device adoption probabilities. For the remainder of the discussion, I refer to the first three segments as “avid readers” and the fourth segment as “general readers.” The estimated genre fixed effects imply that an avid reader buys 8.9 more books than a general reader on average every year. The coefficients on observed income and age imply that older consumers enjoy reading more and that consumers in higher income groups have lower price elasticities of Kindles and books.

The estimates on the e-format tastes show that consumers’ format preferences and substitution patterns between paperbacks and e-books vary across book genres. The estimated e-format genre fixed effects show that if there were no price differences between the two formats, the same consumer would buy 0.92 more “lifestyle” books, 2.76 more “casual” books, and 1.21 fewer “practical” books in e-format in comparison with his paperback consumption. It means that consumers enjoy extra utilities from reading “lifestyle” and “casual” books in e-format and face disutilities from reading “practical” books in e-format.³⁷ Consequently, their substitution patterns between formats differ across genres. I find that “casual” e-books have a lower own-elasticity and a higher cross-elasticity with respect to paperbacks, indicating that they are stronger substitutes for paperbacks. For other coefficient estimates, I find that older consumers dislike the e-format and that e-book variety positively affects e-format attractiveness.

Finally, the paperback retailer choice estimates show that consumers are migrating from offline to online, and from other online retailers to Amazon.com. One

³⁷A caveat is that I cannot rule out the possibility that the e-format fixed effects partially capture different e-book availability across genres. The number of e-books available in the data is aggregate and not genre-specific. If the e-book format is more available in one genre than in another, it can be absorbed by the e-format genre fixed effects. Still, the fixed effects could capture genre differences for utility-related reasons such as different requirements on image display and depth of reading.

interesting finding is that there is positive correlation between Kindle ownership and Amazon retailer choice. I find that a Kindle owner has 54.8% probability of buying paperbacks from Amazon conditional on buying, while the probability is 39.4% for a Kindle nonowner. This could be either a spillover effect (i.e., state dependence) or brand loyalty (i.e., unobserved persistent preference for Amazon).

Consumer heterogeneity. Consumers in the model are heterogeneous in their unobserved baseline reading tastes θ_{ig} and observed demographics such as income and age. The estimates show that the unobserved heterogeneity leads to a much larger difference in demand elasticities and consumption patterns than the observed heterogeneity does. I thus focus on the distinction between avid readers and general readers for the rest of the analysis.

Table 4 compares the demand elasticities and consumption behaviors of a typical avid reader and a typical general reader given average observed prices.³⁸ I find that avid readers have higher probabilities of adopting Kindles and buying them earlier. Their predicted Kindle penetration rate (38%) was higher than general readers' (2.2%) by the end of 2012. I also find that both avid readers and general readers buy more books after adopting Kindles. A typical avid reader buys 12.16 paperbacks per year before adopting a Kindle and buys 3.31 paperbacks and 14.95 e-books once he adopts the e-reader. A typical general reader buys 1.01 paperbacks before adopting a Kindle and buys 0.66 paperbacks and 4.30 e-books once he adopts the e-reader.

The key demand-side finding that drives the supply-side pricing strategy is the relative demand elasticities between Kindles and books for avid and general readers. I find that avid readers have lower price elasticities for both Kindles and e-books than general readers in absolute terms, yet they are relatively more price elastic to e-book prices than to Kindle prices. General readers are more price elastic to Kindle prices

³⁸The demand elasticities of books I find are comparable to extant literature on book consumption (e.g., Chevalier and Goolsbee, 2003; Ghose, Smith, and Telang, 2006; De los Santos, Hortaçsu, and Wildenbeest, 2012; Reimers and Waldfogel, 2014).

Table 4: Consumer Heterogeneity in Kindle and Book Purchases

	Avid reader	General reader
Segment size	7.5%	92.5%
Demand elasticity		
Kindle	-1.41 (0.722)	-6.01 (1.711)
E-book	-0.40 (0.003)	-0.77 (0.006)
Ratio: Kindle/E-book	3.53	7.81
Book consumption per person per year		
Kindle nonowner: # paperbacks q^{P0}	12.16 (0.002)	1.01 (0.004)
Kindle owner: # paperbacks q^{P1}	3.31 (0.057)	0.66 (0.007)
Kindle owner: # e-books q^E	14.95 (0.138)	4.30 (0.078)
Gain from converting a nonowner to an owner		
for the industry $(p^P q^{P1} + p^E q^E) - p^P q^{P0}$	\$-10.98 (0.345)	\$35.69 (0.652)
for Amazon $[(p^P - w^P) q^{P1A} + 0.3p^E q^E] - (p^P - w^P) q^{P0A}$	\$35.67 (0.321)	\$12.46 (0.219)
Cannibalization: # books $q^{P0} - q^{P1}$	8.85 (0.592)	0.35 (0.006)
Cannibalization rate (%): $(q^{P0} - q^{P1}) / q^E$	59.2 (0.168)	8.1 (0.043)

Notes: The values are calculated using the average observed prices across years. Standard errors are in parentheses. The e-book price is $p^E = \$9.72$, the paperback price is $p^P = \$17.66$, and the wholesale price is $w^P = \$15$. The gain from converting a nonowner to an owner is calculated for the analysis in Chapter 7.

than to e-book prices. The intuition is that avid readers buy more books and spend more on books than on devices relative to general readers. Consequently, they care more about subsequent book purchases when considering buying Kindles.

Note that avid readers and general readers differ *only* in their baseline reading tastes θ_{ig} . They do not differ in the price coefficient. The difference in θ_{ig} drives their difference in book usage intensity, device adoption probability, relative demand elasticity between the two products, and in turn optimal dynamic pricing strategies for the firm. The difference in avid readers' and general readers' relative demand elasticities is a result of their endogenous choices rather than being directly specified by the model.

CHAPTER 6

Supply-Side Policy Simulations

Given the estimated demand system, I numerically solve for Amazon's optimal IPD policies with complementary products. I begin by discussing how the existence of the complementary product affects the IPD of the primary product. I hold the complementary product price flat and solve for the optimal IPD strategy for the primary product only. This is comparable to the traditional single-product IPD case and helps illustrate the novelty of the complementary product setting. I then allow the firm to conduct IPD on both the primary and the complementary products. I show that the firm can benefit from a joint IPD policy by exploiting a new dimension of consumer heterogeneity.

There are two note-worthy points. First, I define harvesting and investing based on mark-ups because cost can drop and drive down prices. Firms adopt a harvesting (investing) strategy if the mark-up decreases (increases) over time. Notice that prices can still drop when firms adopt an investing strategy if costs drop faster than prices. Second, the book price I solve is the average price level across book titles. In practice, firms can induce the change in average price through adjusting individual book title prices.

6.1. IPD on the Primary Product Only

In this section, I solve for the optimal IPD policy on the primary product, holding the complementary product price exogenously fixed and flat. Compared with the traditional IPD literature, some results are consistent and some are novel in this setting. In particular, the firm loses more from consumers' strategic behaviors, although complementarity can enhance the firm's IPD ability on the primary product.

First, complementarity provides extra investing incentives to the firm. The optimal pricing policy is investing rather than harvesting as in the traditional IPD literature. Given the observed e-book price of \$9.72, the optimal Kindle price drops from \$371 in 2008 to \$249 in 2012. The markup increases from \$135 to \$160, suggesting an investing strategy. The incentive to penetrate the market early outweighs the incentive to skim high-valuation consumers.

Second, complementarity influences the firm's IPD ability on the primary product by changing the demand elasticities of the primary product. In particular, a lower e-book price increases the difference in the demand elasticities of Kindles between the two consumer types. This enhances the firm's IPD ability so that the firm can invest less in Kindles. When the e-book price decreases from \$9.94 to \$9.63, the optimal Kindle price path shifts up, and the average price increases from \$292 to \$305.

Third, it is worth noting that the profit loss from consumers' forward-looking behavior is higher in the complementary product setting. Traditional single-product IPD literatures conclude that profits and prices are higher if consumers are myopic and do not intertemporally arbitrage. For instance, Nair (2007) studies IPD for video games and finds that the profit under myopic consumers is 172.2% higher than that under forward-looking consumers. This number is 284% in my case because the difference in profits comes from not only the primary product sales but also from the subsequent complementary product sales. The optimal price is higher under myopic

consumers. It drops from \$470 to \$358 if consumers are myopic and drops from \$371 to \$249 if consumers are forward-looking.

6.2. IPD on Both Products

In this section, I solve for the optimal joint IPD policies on both products (hereon, the joint IPD case).


I find that the shapes of the optimal pricing functions differ for avid readers and general readers. This result is driven by the heterogeneity in relative demand elasticities, which is novel in the complementary product setting. Optimal policies are functions of the penetration rates of consumer types as shown in Figure 5. As the penetration rate of avid readers increases, the optimal strategy is to harvest on Kindles and invest in e-books. The opposite is true for general readers.³⁹ In general, it is optimal to invest in the product with higher relative demand elasticities and harvest on the product with lower relative demand elasticities.⁴⁰ Avid readers are more price elastic to e-books than to Kindles. The opposite is true for general readers. The joint IPD policy exploits this new dimension of consumer heterogeneity. The overall price path balances the incentives for both consumer types and depends on the mix of consumers in the market.

I use the optimal pricing functions to simulate the price trajectories and market outcomes. The predicted price trajectory is to harvest on Kindles and invest in e-books as shown in Table 5 under specification (i). The predicted e-book price increases

³⁹Notice that consumer types are still unobserved to the firm. This result is a characteristic of the policy functions, which are functions of the number of consumers in each type.

⁴⁰This result is comparable to the static third-degree price discrimination problem, where consumer types are observable (e.g., Aguirre et al., 2010). Compared with the nondiscriminatory case, prices are lower in the weak market, where the demand is more price sensitive, and higher in the strong market. I study IPD when consumer types are unobserved. The results can still be qualitatively applied to the scenario with observed consumer types (e.g., firms know consumer types from transaction history or surveys). Based on the results, the firm should provide promotions on Kindles to general readers and promotions on e-books to avid readers. Dynamic pricing is no longer necessary in this case; the firm can charge a flat optimal price for each type.

Figure 5: Policy Functions: Investing vs. Harvesting



by 11.1% over five years. To implement the e-book price increase in practice, Amazon can frame the low initial prices as introductory prices. I compare the results to a single IPD case in which the firm conducts IPD only on Kindles and optimally commits to a fixed e-book price.⁴¹ The predicted price trajectory in the single IPD case is to invest in Kindles. Amazon should start with a relatively smaller mark-up for Kindles and gradually increase it over time. The retail price still drops over time as the cost drops faster. The joint IPD policy induces faster-declining Kindle prices and lower penetration rates of both avid readers and general readers.

The joint IPD policy, as compared with the single IPD policy, benefits the firm in two ways. First, it offers the firm a better screening device and induces a higher fraction of avid readers, who are more profitable, to adopt Kindles as shown in Figure 6. To see how the screening device works, consider two scenarios. With only one product, raising the price will discourage both avid readers and general readers from buying; they respond in the same direction. With two products, raising the Kindle price and reducing the e-book price properly will attract avid readers and discourage


⁴¹I do not examine how Amazon gets commitment power. In practice, commitment power can come from reputation (e.g., Apple's strategy on hardware) or contractual arrangement (e.g., resale price maintenance).

Table 5: IPD on Both Products vs. IPD on the Primary Product Only

	(i)		(ii)	
	Single IPD	Joint IPD	Single IPD	Joint IPD
Kindle price path (\$, markup in parentheses)				
2008	318 (82)	347 (111)	215 (-21)	328 (92)
2009	278 (93)	268 (83)	187 (2)	267 (82)
2010	247 (103)	218 (74)	167 (23)	212 (68)
2011	224 (111)	184 (71)	155 (42)	173 (60)
2012	207 (118)	159 (60)	148 (59)	146 (57)
E-book price path (\$)				
2008	9.66	9.79	10.70	9.87
2009	9.66	10.29	10.70	10.54
2010	9.66	10.66	10.70	10.93
2011	9.66	10.87	10.70	11.24
2012	9.66	11.01	10.70	11.48
Penetration rate by 2012 (%)				
avid readers	38.5%	27.6%	33.9%	27.1%
general readers	2.87%	2.36%	3.07%	2.30%
Discounted profits (2008-2012, million \$)				
total	10,576	9,765	22,190	23,027
from Kindles	1,796	1,000	956	2,200
from books	8,780	8,765	21,233	20,827

Notes: Specification (i) shows the scenario in which the share of avid readers in the initial market is at the estimated level of 7.47%. Specification (ii) shows the scenario in which the share is 30%.


Figure 6: Fractions of Avid Readers in Kindle Sales



general readers, while the opposite price change will discourage avid readers and attract general readers. In other words, manipulating two price instruments can induce different demand responses from avid and general readers. The firm can use different price path combinations to induce different consumer types to make purchases. This echoes the importance of modeling self-selection based on consumer heterogeneity.

Second, the joint IPD policy limits consumers' strategic behaviors by providing conflicting incentives to arbitrage on price changes over time. Instead of delaying purchase as in the traditional literature, consumers adopt Kindles earlier even at higher Kindle prices. The firm not only extracts more Kindle profits but also earns more book profits from subsequent book sales. Figure 7 plots the Kindle sales by consumer type over time under the single IPD case and the joint IPD case. The total Kindle sales are normalized to 100 for each type in each case so that I can focus on the consumer arbitrage behavior. In the single IPD case, avid readers delay purchase as the Kindle price drops from \$318 in 2008 to \$207 in 2012 (markup increases, so still "investing"). Interestingly, they adopt Kindles earlier in the joint IPD case, even though the Kindle price is higher and drops faster from \$347 to \$159. This is driven by increasing e-book prices over time. General readers are much less responsive. In the

Figure 7: Kindle Sales Over Time



Notes: To focus on the sales distribution over time, I normalize the total sales of each consumer type to 100 under each scenario.

traditional IPD literature, forward-looking consumers strategically delay purchase and hurt firm profitability. In the joint IPD case, consumers have two conflicting incentives: They should delay purchase given declining Kindle prices but should adopt earlier given increasing e-book prices. This limits their ability to intertemporally arbitrage.

The two mechanisms discussed above contribute to the profitability of jointly conducting IPD on Kindles and e-books. However, profitability is not guaranteed. Given that consumers will arbitrage on price changes, it might be better to commit to a fixed price. This is also true in the traditional single-product IPD. I find that profitability of the joint IPD increases when there are more avid readers. Table 5 compares the scenario in which the share of avid readers in the initial market is at the estimated level of 7.47% (specification i) with the scenario in which the share is 30% (specification ii). If the fraction of avid readers is 30%, Amazon will benefit substantially from the joint IPD policy; given that the observed Kindle ecosystem and its related business produced \$10.6 billion profits from 2008 to 2012, the total profit


gain is \$837 million.⁴² If the fraction of avid readers is too low as estimated from the data, it is better to commit to a fixed e-book price. Committing to a fixed e-book price is better under the estimated market composition for two reasons. First, under the joint IPD policy, the firm can harvest avid readers using high Kindle prices while keeping Kindles attractive using low e-book prices. Yet low e-book prices cannot effectively keep general readers, who are less price sensitive to books; high Kindle prices substantially discourage them from adopting. The firm needs a large enough share of avid readers in the initial market to induce profitability. Second, the degree of harvesting/investing depends on the mix of consumer types. If the fraction of general readers is too high, the firm would adjust the price level to accommodate the general readers, resulting in a deviation from the optimal pricing strategy to attract avid readers.

6.3. Discussion

The last section presents the optimal joint IPD policy and discusses its advantages and profitability. In this section, I compare Amazon's observed pricing strategy with the proposed policy. Given that the estimated avid reader share is 7.47%, the model predicts that Amazon should commit to a fixed e-book price and invest in Kindles by starting with a low mark-up. Figure 8 shows that in practice Amazon also committed to a fixed e-book price but reduced Kindle prices faster than what the model would suggest. Note that the focus of this dissertation is not to rationalize the observed policy. The fact that the observed policy differs from the proposed policy does not imply that Amazon behaves suboptimally. The difference arises because in practice Amazon is solving for a more complex problem. It faced competition from Barnes & Noble's NOOK and Apple's iPad after 2010, which provided pricing-cutting incentives

⁴²See <http://allthingsd.com/20130812/amazon-to-sell-4-5-billion-worth-of-kindles-this-year-morgan-stanley-says/> and Amazon Annual Financial Statements 2010, 2012.

Figure 8: Observed Prices vs. Simulated Prices



to maintain market share. It lost the e-book pricing right from 2010 to 2012 because of an e-book contract switch.⁴³ It also made innovation choices along with the pricing choices. I do not have data on competitors and R&D. Incorporating all these factors is also computationally prohibitive. Yet the monopolist’s dynamic pricing problem in a complementary product setting is both important and interesting to address. I provide a framework to understand the basic trade-offs and how firms can possibly benefit from joint IPD strategies in this novel setting.

The results shed light on the pricing incentives behind Amazon’s reputation for “pricing-at-cost” for Kindles and e-books. I find some evidence that Amazon chooses this strategy because of business spillover effects from e-books to other products. First, there is a positive correlation between Kindle ownership and buying paperbacks from Amazon. If the correlation comes from a spillover effect, it means that e-reading brings additional paperback sales to Amazon. Second, there is a spillover effect to other product categories. In the simulation, I allow for a spillover effect per book transaction to obtain a more realistic marginal cost of selling books for Amazon. The

⁴³The model can be extended to accommodate the contract change by modifying Amazon’s profit function. The model-predicted prices are comparable to the observed Kindle prices during the new contract period, providing another validation of the model setup.

calibrated spillover effect is \$5.50 per book transaction.⁴⁴ Survey results also exhibit a similar spillover effect: Kindle owners spend 56% more with Amazon as compared with Kindle nonowners (Consumer Intelligence Research Partners, 2012).⁴⁵ As the Kindle ecosystem comprised 11% of Amazon's total revenue in 2013, the existence of this spillover effect could provide large incentives for Amazon to promote the e-book business.⁴⁶

⁴⁴To validate this number, notice that Kindles with no ads are \$30 more expensive than those with ads. A typical Kindle owner buys 5.2 more books than a nonowner and produces \$28.6 extra spillover effect based on the model prediction. Kindle ad revenue is comparable to one-year additional spillover effect due to Kindle adoption. Model predictions are robust to other spillover effect sizes.

⁴⁵See <http://www.publishersweekly.com/pw/by-topic/digital/devices/article/50601-the-kindle-multiplier.html>.

⁴⁶See <http://allthingsd.com/20130812/amazon-to-sell-4-5-billion-worth-of-kindles-this-year-morgan-stanley-says/>.

CHAPTER 7

Cannibalization or Market Expansion? The Impact of E-Books on Print Book Sales

I can use the estimated demand system to explore the impact of e-books on print book sales. The model is well suited to address this issue because it captures the substitution pattern between paperbacks and e-books and the device adoption decision. The former determines the individual-level cannibalization and market expansion effects, and the latter influences the number of consumers who read e-books and generate the two effects. The components of the demand model are directly linked to the cannibalization and market expansion effects. Cannibalization happens when two conditions hold: (1) the consumer would buy paperbacks in the absence of e-books, and (2) the consumer would prefer e-books to paperbacks when he can choose from both formats. Market expansion happens when the second condition holds and the first condition does not hold. In the model, the first condition is determined by the baseline taste parameters θ_{ig} . The second condition is determined by the e-format taste parameters θ_g^E and the price differences between e-books and paperbacks.

In this chapter, I explore the degrees of cannibalization and market expansion effects given observed e-reader and e-book prices. I further ask how the effect sizes would change if e-book prices increase by \$2. The motivation of the counterfactual comes from a policy change in practice. Publishers were concerned about the cannibalization effect of the low-priced e-books. They sought to raise e-book prices of the new releases and *New York Times* best sellers by \$2 in 2010. I ask whether cannibalization is indeed weaker and whether the publishing industry would benefit from this price change. Amazon was deprived of the e-book pricing right and collected 30% of the book revenues from 2010 to 2012. I construct the firm problem similarly as in Section 3.3 except that Amazon only solves for optimal Kindle prices in this chapter. The book profits generated by each Kindle owner becomes $r^1 = \lambda p^E q^E + (p^P - w^P) q^{P1A}$, where $\lambda = 0.3$ is the fraction of e-book revenue shared by Amazon.

7.1. Effect Sizes Under Observed Pricing

I first examine the impact of e-books given current e-book and e-reader prices. At the individual level, I use consumers' book consumption before and after Kindle adoption to evaluate the gain (loss) from converting a Kindle nonowner to an owner. At the aggregate level, I quantify the degrees of cannibalization and market expansion effects and evaluate the gain (loss) from introducing e-books. The results illustrate how e-books affect consumers' book consumption and how beneficial it is for publishers and Amazon to promote e-reading.

The calculation proceeds as follows: Given the demand estimates and the observed Kindle and e-book pricing, I simulate the market without e-books and calculate the paperback sales S^{P0} that would have occurred on Amazon, on other online retailers, and at offline bookstores. I then take the difference of the simulated sales S^{P0} and observed sales S^{P1} to get the degree of cannibalization and market expansion

effects, overall and for each retailer. In particular, I denote the total e-book sales as S^E . Part of the e-book sales comes from cannibalization, which is represented by the loss of paperback unit sales due to e-books $S^{P0} - S^{P1}$. The rest of the e-book sales comes from market expansion, which is represented by the incremental book sales in e-format that would not otherwise have occurred in paperback format $S^E - (S^{P0} - S^{P1})$.⁴⁷ The cannibalization rate is defined as the percentage of total e-book sales that comes from cannibalization $\frac{S^{P0} - S^{P1}}{S^E}$. In terms of book revenue, cannibalization loss is defined as the number of cannibalization e-books times the difference in average paperback and e-book prices $(S^{P0} - S^{P1})(p^P - p^E)$. Market expansion gain is defined as the number of market expansion e-books times the average e-book price $[S^E - (S^{P0} - S^{P1})]p^E$. The difference between cannibalization loss and market expansion gain is the net impact of e-books on the industry profits. Individual-level calculation follows the same procedure except that I substitute individual book quantity choices for aggregate book sales.

At the individual consumer level, I find that both avid readers and general readers buy more books after adopting Kindles as shown in Table 4. Yet converting them to Kindle owners (i.e., e-book readers) is not always profitable. Besides, the conversion affects publishers and Amazon differently. General readers have lower industrywide cannibalization rates (8.1%), and converting them is more profitable for the industry. Avid readers have higher industrywide cannibalization rates (59.2%) and are easier to convert, while converting them is more profitable for Amazon.⁴⁸ In terms of dollars,

⁴⁷This market expansion effect does not include incremental paperback sales ΔS^P caused by e-books through, for instance, word-of-mouth effect from e-book consumers. If such an effect exists, my approach would underestimate the cannibalization effect by ΔS^P because $S^{P1} = S^{P0} + \Delta S^P - \text{cannibalization}$. The model does not contain a network effect to capture ΔS^P . However, I believe that word-of-mouth is more relevant at the book title level, while my model is at the genre level. Still, I am able to capture consumers' incremental book consumption once they start e-reading.

⁴⁸Avid readers and general readers have different cannibalization rates because two conditions jointly induce cannibalization: 1) the consumer will buy paperbacks in the absence of e-books, and 2) the consumer prefers e-books to paperbacks when he can choose from both formats. The two types have the same probability of satisfying the second condition because they share the same e-format preference. The avid readers, however, are much more likely to satisfy the first condition, as they have higher reading tastes.


converting a typical avid (general) reader to a Kindle owner leads to a \$10.98 loss (\$35.69 gain) per year for the industry. In contrast, converting a typical avid (general) reader leads to a \$35.67 gain (\$12.46 gain) per year for Amazon. The difference comes from two sources: (1) Amazon is only one of the paperback retailers that the industry cares about, and (2) Amazon earns additional paperback profits from Kindle owners who are more likely to choose Amazon to buy paperbacks. The difference may provide Amazon and publishers with different incentives to promote e-reading. For instance, converting avid readers leads to a loss for the industry. However, source (2) is large enough for Amazon to benefit from it and seek to attract avid readers even at the expense of the industrywide profits.

Accounting for the size and mix of Kindle owners yields the impact of e-books at the aggregate level. I find that the industry benefits from the introduction of e-books because the overall market expansion gain is larger than the overall cannibalization loss. Yet the cannibalization loss is unevenly born by different paperback retailers. In Figure 9, I plot the cannibalization and market expansion effects and the decomposition of cannibalization effect by retailer over time. I find that 42% of e-book sales come from cannibalizing paperbacks and that 58% come from market expansion.⁴⁹ Of the cannibalization effect, offline bookstores bear 53% of the loss, other online retailers bear 15%, and Amazon bears 32%.

The simulation results also shed light on how the effect sizes would change in the future. I find that the overall cannibalization rate drops from 45.7% in 2008 to 39.6% in 2012 and that the cannibalization burden is switching to Amazon over time. This

⁴⁹The results can be validated using model-free evidence from the data. There is a limited set of households whose book transactions are observed before and after Kindle adoption. Comparing their book consumption as a Kindle nonowner and owner, I can obtain a cannibalization rate of 40.5%. This calculation does not account for heterogeneous consumer tastes. There is also qualitative evidence from the industry. Lulu.com is a leading print and digital self-publishing service provider for over 1 million authors. It reported that those authors who published their books in both print and e-book formats tended to sell double the amount of books (<http://www.the-digital-reader.com/2012/04/20/ebook-sales-dont-undercut-print-sales-lulu-reports/\#.U3IjQa1dU00>). The cannibalization effect is much smaller than what traditional publishers thought it to be.

Figure 9: Cannibalization and Market Expansion



is likely to continue given the current industry conditions. There are two underlying driving forces for the trend. First, the mix of Kindle owners evolves over time. Avid readers are the early adopters. As their market saturates, general readers start to constitute a larger share of Kindle owners, and they have lower cannibalization rates. It means that the industry would benefit even more from e-books in the future. Second, consumers are migrating from offline to online, and from other retailers to Amazon. As Amazon becomes a larger paperback retailer over time, it shoulders more cannibalization burden. Amazon can still benefit from the trend as long as the increase in cannibalization burden is smaller than the increase in book sales.

7.2. Effect Sizes Under Alternative E-Book Prices

The previous section explores the impact of e-books under observed prices. Publishers and retailers can change e-reader and e-books prices to change the magnitudes of the

cannibalization and market expansion effects. In particular, e-book prices affect the individual-level effect, or the gain (loss) from an additional Kindle owner. E-book and Kindle prices jointly affect the aggregate-level effect by determining the size and mix of Kindle owners.

In this section, I use a simulation to explore whether raising e-book prices by \$2, similar to what publishers did in 2010, would reduce cannibalization and benefit the publishing industry. Accounting for Amazon’s Kindle pricing response is important to obtain the full effect of this e-book price change. I first solve for Amazon’s optimal dynamic pricing strategies given the observed and new e-book prices. I then simulate the Kindle price paths and compare the market outcomes under three scenarios: (1) the e-book price is at the observed level of \$9.72; (2) the e-book price increases to \$11.72, while the Kindle prices remain unchanged; and (3) the e-book price increases to \$11.72, and Amazon responds by adjusting the Kindle prices optimally. The results are presented in Table 6.

I first keep Kindle prices unchanged as in Column (ii). At the individual level, a higher e-book price reduces e-book sales per Kindle owner as well as the size of cannibalization and market expansion each owner creates (measured in number of books). A higher e-book price also makes Kindle less attractive and reduces the number of Kindle owners. The aggregate effect is that both cannibalization loss and market expansion gain drop. The net effect on the industry profits is negative: The benefit from e-book introduction decreases substantially by 67.5%.

When allowing Kindle prices to change as in Column (iii), I find that the situation is even worse. Interestingly, a higher e-book price induces higher Kindle prices, according to the firm’s first-order condition. This is because consumers buy fewer e-books and provide lower gains for Amazon to convert nonowners to owners. The higher Kindle prices further reduce the number of Kindle owners. Cannibalization loss and market expansion gain drop even more, and the industry’s benefit from e-book

Table 6: Effects of Raising E-book Prices

	(i)	(ii)	(iii)
Kindle Price Path (\$)			
2008	174	174	292
2009	141	141	247
2010	116	116	212
2011	99	99	186
2012	88	88	167
Penetration rate by 2012 (%)			
avid readers	34.1%	29.0%	13.0%
general readers	4.9%	4.6%	2.6%
Industry revenues (2008-2012, \$)			
cannibalization loss	36,857	10,300	5,082
market expansion gain	44,148	12,670	6,589
net gain	7,291	2,370 (-67.5%)	1,507 (-79.3%)

Notes: In specification (i), the e-book price is \$9.72; in specification (ii), the e-book price increases to \$11.72, while Kindle prices remain unchanged; in specification (iii), the e-book price increases to \$11.72, and Kindle prices are allowed to adjust optimally. The magnitudes of the revenues are for the data sample. For the last row, percentage changes from specification (i) are in parentheses.

introduction drops by 79.3%.

The simulation results imply that increasing e-book prices by \$2 makes the industry worse off. The key is that higher e-book prices induce higher Kindle prices, which substantially discourage general readers from adopting Kindles. The loss from a smaller market expansion effect outweighs the gain from a smaller cannibalization effect. The simulation also illustrates the misaligned incentives between publishers and Amazon. Amazon and publishers differ in their gains from e-books. While Amazon adjusts Kindle prices optimally for itself, publishers can get hurt in the process. Kindle serves as the gateway product to e-reading, yet its pricing may work against the diffusion of e-reading.

CHAPTER 8

Conclusion

I estimate a dynamic discrete-continuous model of consumer e-reader and book purchases and use the estimated demand system to simulate firm strategies and market outcomes. The results provide a better understanding of consumer demand and firm strategy in the e-book market.

Chapter 2 to Chapter 6 provides a framework to understand the joint IPD strategy for hardware and software. I start with estimating a dynamic demand model of e-readers and e-books using individual transaction data. I then use the estimated demand system to compute the firm's optimal dynamic pricing strategy. Two modeling decisions are important: (1) accounting for consumer heterogeneity in reading tastes, as different types of consumers generate different book revenues and respond differently to price changes; and (2) accounting for the dynamic device adoption decision, as it allows for self-selection based on heterogeneous consumer tastes.

The demand-side estimation reveals a novel dimension of consumer heterogeneity that firms can exploit in the complementary product setting. The supply-side simulation proposes a novel joint IPD strategy from which firms can benefit in two ways. Firms can better screen consumers and limit consumers' ability to intertemporally arbitrage. The profitability of the joint IPD strategy increases as the percentage of avid readers increases in the initial market. The study provides new insights into the

traditional pricing approach for books. Traditional book prices are set at standard discounts off list prices and remain flat over time. I find that retailers can conduct IPD on e-books to improve profitability, especially given that e-book prices are easily adjustable because of their digital nature.

The key underlying mechanism of the joint IPD policy is the following. The usage intensity of the software drives the adoption of the hardware. It leads to the difference in the relative demand elasticity between hardware and software. This demand heterogeneity further leads to the difference in supply-side optimal price trajectory combinations of hardware and software. The results are applicable to other industries such as consoles and video games, Apple TVs and digital content on iTunes, razors and blades, printers and cartridges, and K-cups and espresso machines. One needs to estimate the relative demand elasticities for different consumer types and the consumer mix. The overall price trajectory and profitability depend on the mix of consumer types.

Chapter 7 examines the impact of e-books on print book sales. At the individual level, both avid readers and general readers buy more books after becoming Kindle owners. Yet a higher percentage of avid readers' e-book consumption comes from cannibalizing paperbacks. General readers have lower cannibalization rates, but are more difficult to convert. Amazon benefits more from converting avid readers and the industry benefits more from converting general readers. The aggregate effect of e-books depends on (1) the individual-level effects, or the gain from converting a Kindle nonowner to an owner, and (2) the cumulative number of Kindle owners and nonowners. The first factor depends solely on e-book prices. The second factor depends on both e-book and Kindle prices. I find that given the observed prices, 42% of the e-book sales come at the expense of paperback sales, while 58% would not have occurred in the absence of e-books. The gain from market expansion exceeds the loss from cannibalization, indicating that e-book is a good opportunity to expand

the book business. I further show that raising e-book prices by \$2 reduces cannibalization loss as well as market expansion gain. The overall effect is negative for the industry. Higher e-book prices might not help eliminate publishers' concerns about cannibalization.

This dissertation has managerial implications for publishers, retailers, and policymakers in the publishing industry. For offline bookstores, there has been a decline in the numbers and sales of their products even before the introduction of e-books. The offline book store sales peaked in 2007, one year before the introduction of Amazon Kindle, and then steadily declined (U.S. Census Monthly Retail Trade Survey). One of the largest chain retailers, Borders, closed down in 2011. The number of booksellers fell to 7,244 establishments by 2012, a reduction of 27.2% from 2007 (Gilbert, 2014). This dissertation addresses how the introduction of the e-book further drove book buyers from offline to online. Offline bookstores have to bear the loss from additional print book sales drop due to e-books (i.e., cannibalization), yet they cannot enjoy the benefit from the additional book sales created by e-books (i.e., market expansion).

Although Amazon also bears the cannibalization loss as a print book retailer, it might benefit from e-books in three ways: (1) additional e-book sales because of market expansion, (2) additional print book sales driven by Kindle adopters (assuming that the positive correlation between Kindle ownership and Amazon retailer choice comes from state dependence), and (3) the spillover effect to other product categories. The second effect suggests that e-books accelerate the shift of book sales from other retailers to Amazon. If the gain from e-books is large enough, Amazon would have strong incentives to promote e-book and e-reader sales, even at a loss. The downside is that Amazon needs to bear a larger cannibalization loss as it becomes a larger paperback retailer.

Publishers have misaligned incentives with Amazon. There are several reasons and consequences. First, publishers care about all retailers, not just Amazon. I find

that converting a general reader is more profitable for publishers, while converting an avid reader is more profitable for Amazon. The difference in gain from converting consumers to e-book readers implies that publishers and Amazon differ in their incentives to promote e-reading. Second, publishers care about book revenue, while Amazon cares about both book and Kindle profits. For example, simulation results show that Kindle prices increase in response to a \$2 increase in e-book prices, suggesting that Amazon's e-reader strategy might not help the diffusion of e-reading. Third, e-books can drive additional paperback sales to Amazon, which provides Amazon with additional incentives to promote e-reading. In particular, the industrywide print book sales can drop, while Amazon's print book sales still increase. In all, publishers need to be alerted about these differences when making e-book related decisions.

Policymakers should take a long-term view on the impact of e-books. Avid readers start e-reading earlier and create more cannibalization in the short run. Yet as more general readers start e-reading, the industrywide cannibalization rate drops. E-books are not just a challenge, but they are also an opportunity. Policymakers should also be aware that publishers and Amazon might have different, even conflicting, incentives and strategies of developing e-reading.

There are possible avenues for future research. First, publishers conspired with Apple to raise e-book prices in 2010 because they worried about both cannibalization and that Amazon was becoming so strong that it would dictate the terms of book sales. An interesting but challenging question is how retailer competition affects cannibalization, market expansion, and e-book and e-reader pricing strategies. My monopoly model illustrates the fundamental trade-offs for further competitive analysis. It is challenging to solve for a dynamic competition model with both e-readers and e-books, yet the multiproduct setting can lead to potentially interesting competition patterns. Second, one could model innovation and quality choices in addition to pricing. These factors may become more important as the e-reader market matures

and more sales come from upgrading. Third, paperback prices do not change in response to e-book introduction both in the model and in practice. Would publishers or paperback retailers benefit from changing paperback prices? How should they change the prices? Answers to these questions can be valuable to the future of the publishing industry in the e-book era.

APPENDIX

A. Likelihood Function Construction

Consumers' optimal format-quantity choices are (summarized from Table 1, subscript t dropped)

$$\{q_{ig}^{P1*}, q_{ig}^{E*}\}_{g=1,2,3} = \begin{cases} \{0, 0\} & \text{if } p_g^P > \frac{a_{ig}^P}{b_i}, p_g^E > \frac{a_{ig}^E}{b_i} \\ \{a_{ig}^P - b_i p_g^P, 0\} & \text{if } p_g^P < \frac{a_{ig}^P}{b_i}, p_g^E > \frac{a_{ig}^E}{b_i}, \text{ or} \\ & p_g^P < \frac{a_{ig}^P}{b_i}, p_g^E < \frac{a_{ig}^E}{b_i}, a_{ig}^P - b_i p_g^P > a_{ig}^E - b_i p_g^E \\ \{0, a_{ig}^E - b_i p_g^E\} & \text{if } p_g^P > \frac{a_{ig}^P}{b_i}, p_g^E < \frac{a_{ig}^E}{b_i}, \text{ or} \\ & p_g^P < \frac{a_{ig}^P}{b_i}, p_g^E < \frac{a_{ig}^E}{b_i}, a_{ig}^P - b_i p_g^P < a_{ig}^E - b_i p_g^E \end{cases}$$

$$q_{ig}^{P0*} = \begin{cases} a_{ig}^P - b_i p_g^P & \text{if } p_g^P < \frac{a_{ig}^P}{b_i} \\ 0 & \text{otherwise} \end{cases}$$

where the taste parameters are parameterized as

$$\begin{aligned} a_{ig}^P &= \bar{a}_{ig}^P + \eta_{ig}^P = \theta_{ig} + \beta_1 D_i^{age} + \eta_{ig}^P \\ a_{ig}^E &= \bar{a}_{ig}^E + \eta_{ig}^E = \underbrace{\theta_{ig} + \beta_1 D_i^{age}}_{\text{Baseline taste}} + \underbrace{(\theta_g^E + \beta_2 D_i^{age} + \beta_3 \log n^E) \cdot 1\{ebook\}}_{\text{E-format taste}} + \eta_{ig}^E \end{aligned}$$

The format-quantity choice probability comes from the error terms in the taste parameters. Define the realized error terms given the quantity choice q_{ig}^P as $\eta(q_{ig}^P) \equiv q_{ig}^P + b p_g^P - \bar{a}_{ig}^P$ and $\eta(q_{ig}^E) \equiv q_{ig}^E + b p_g^E - \bar{a}_{ig}^E$. Define the thresholds of worth buying as $\bar{\eta}_{ig}^P \equiv b p_g^P - \bar{a}_{ig}^P$ and $\bar{\eta}_{ig}^E \equiv b p_g^E - \bar{a}_{ig}^E$. To simplify the notation, I drop i and g subscripts for now.

Case 1: A Kindle nonowner. The optimal quantity choice requires that the

normally distributed error terms satisfy

$$\begin{cases} \eta^P = \eta(q^P) > \bar{\eta}^P & \text{if } q^P > 0 \\ \eta^P \leq \bar{\eta}^P & \text{if } q^P = 0 \end{cases}$$

The probability of buying $q^P > 0$ number of books is

$$f(\eta^P = \eta(q^P) \mid \eta^P > \bar{\eta}^P) \Pr(\eta^P > \bar{\eta}^P) = \frac{1}{\sigma} \phi\left(\frac{\eta(q^P)}{\sigma}\right)$$

The probability of buying $q^P = 0$ is

$$\Pr(\eta^P < \bar{\eta}^P) = \Phi\left(\frac{\bar{\eta}^P}{\sigma}\right)$$

Here $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF of the standard normal distribution. So the contribution to the likelihood function for a Kindle nonowner is

$$l^{nonowner} = \prod_g \left[1 \{q_g^P > 0\} \frac{1}{\sigma} \phi\left(\frac{\eta(q_g^P)}{\sigma}\right) + 1 \{q_g^P = 0\} \Phi\left(\frac{\bar{\eta}_g^P}{\sigma}\right) \right]$$

Case 2: A Kindle owner. Similarly, the optimal quantity choice requires that the error terms satisfy

$$\begin{cases} \eta^P < \bar{\eta}^P, \eta^E < \bar{\eta}^E & \text{if } q^P = 0, q^E = 0 \\ \eta^P = \eta(q^P) > \max\{\bar{\eta}^P, \eta^E + (\bar{\eta}^P - \bar{\eta}^E)\} & \text{if } q^P > 0, q^E = 0 \\ \eta^E = \eta(q^E) > \max\{\bar{\eta}^E, \eta^P - (\bar{\eta}^P - \bar{\eta}^E)\} & \text{if } q^P = 0, q^E > 0 \end{cases}$$

Intuitively, paperbacks are chosen because they are worth buying ($\eta^P > \bar{\eta}^P$) and they are better than e-books ($\eta^P > \eta^E + (\bar{\eta}^P - \bar{\eta}^E)$), vice versa. The probability

that $q^P = 0, q^E = 0$ is $\Pr(\{0, 0\}) = \Pr(\eta^P < \bar{\eta}^P) \Pr(\eta^E < \bar{\eta}^E)$. The probability that $q^P > 0, q^E = 0$ is $\Pr(\{q^P, 0\}) = f(\eta^P = \eta(q^P) \mid \eta^P > \max\{\bar{\eta}^P, \eta^E + (\bar{\eta}^P - \bar{\eta}^E)\}) \cdot \Pr(\eta^P > \max\{\bar{\eta}^P, v^E + (\bar{\eta}^P - \bar{\eta}^E)\})$. Similar for the probability $\Pr(\{0, q^E\})$. The last two probabilities are conditional probabilities of a truncated normal distribution where the truncation point is a result of a maximization operator. I use the quadrature method to calculate it. The details are presented in Appendix B.

B. Probability and Indirect Utility Calculation

Probability Calculation Notations follow Appendix A. I need to calculate $\Pr(\{q^P, 0\})$, or the following density

$$f(\eta^P = \eta(q^P) \mid \eta^P > \max\{\bar{\eta}^P, \eta^E + (\bar{\eta}^P - \bar{\eta}^E)\}) \Pr(\eta^P > \max\{\bar{\eta}^P, \eta^E + (\bar{\eta}^P - \bar{\eta}^E)\})$$

where η^P and η^E are i.i.d. normally distributed error terms with mean 0 and variance σ^2 . $\bar{\eta}^P$, and $\bar{\eta}^E$ are known deterministic parts. Define $\Lambda \equiv (\bar{\eta}^P - \bar{\eta}^E)$ to simplify the discussions below. It is easier to start with calculating the CDF instead of the PDF: $\Pr(\eta^P \leq \eta(q^P) \mid \eta^P > \max\{\bar{\eta}^P, \eta^E + \Lambda\}) \cdot \Pr(\eta^P > \max\{\bar{\eta}^P, \eta^E + \Lambda\})$. Define $\eta^P \leq \eta(q^P)$ as event A , $\eta^P > \bar{\eta}^P$ as event B , $\eta^P > \eta^E + \Lambda$ as event C , and $\bar{\eta}^P > \eta^E + \Lambda$ as event D . Then the CDF can be written as $\Pr(A \mid B \cap C) \Pr(B \cap C)$. Notice that event $B \cap D$ implies event C , and event $C \cap \neg D$ implies event B . Event B and event D are independent. For the first component,

$$\begin{aligned} \Pr(A \mid B \cap C) &= \Pr(A \mid B \cap C \cap D) \Pr(D) + \Pr(A \mid B \cap C \cap \neg D) \Pr(\neg D) \\ &= \Pr(A \mid B \cap D) \Pr(D) + \Pr(A \mid C \cap \neg D) \Pr(\neg D) \\ &= \frac{\Pr(A \cap B \cap D)}{\Pr(B \cap D)} \Pr(D) + \frac{\Pr(A \cap C \cap \neg D)}{\Pr(C \cap \neg D)} \Pr(\neg D) \end{aligned}$$

where $\Pr(B) = \left[1 - \Phi\left(\frac{\bar{\eta}^P}{\sigma}\right)\right]$, $\Pr(D) = \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right)$, $\Pr(\neg D) = 1 - \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right)$, $\Pr(B \cap D) = \Pr(B) \Pr(D)$ and

$$\begin{aligned}
\Pr(A \cap B \cap D) &= \left[\Phi\left(\frac{\eta(q^P)}{\sigma}\right) - \Phi\left(\frac{\bar{\eta}^P}{\sigma}\right) \right] \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right) \\
\Pr(A \cap C \cap \neg D) &= \int_{\bar{\eta}^P+1}^{\eta(q^P)} \left[\Phi\left(\frac{x - \Lambda}{\sigma}\right) - \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right) \right] dF_x \\
&= \int_{\bar{\eta}^P+1}^{\eta(q^P)} \Phi\left(\frac{x - \Lambda}{\sigma}\right) dF_x - \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right) \Phi\left(\frac{\eta(q^P)}{\sigma}\right) \\
\Pr(C \cap \neg D) &= \int_{\bar{\eta}^P}^{+\infty} \left[\Phi\left(\frac{x - \Lambda}{\sigma}\right) - \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right) \right] dF_x \\
&= \int_{\bar{\eta}^P}^{+\infty} \Phi\left(\frac{x - \Lambda}{\sigma}\right) dF_x - \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right) \left[1 - \Phi\left(\frac{\bar{\eta}^P}{\sigma}\right)\right]
\end{aligned}$$

For the second component

$$\begin{aligned}
\Pr(B \cap C) &= \Pr(B \cap C | D) \Pr(D) + \Pr(B \cap C | \neg D) \Pr(\neg D) \\
&= \Pr(B \cap C \cap D) + \Pr(B \cap C \cap \neg D) \\
&= \Pr(B \cap D) + \Pr(C \cap \neg D) \\
&= \Pr(B) \Pr(D) + \Pr(C \cap \neg D) \\
&= \int_{\bar{\eta}^P}^{+\infty} \Phi\left(\frac{x - \Lambda}{\sigma}\right) dF_x
\end{aligned}$$

In all, $CDF(\eta(q^P)) = \Pr(A | B \cap C) \Pr(B \cap C) = \left[\frac{(a-b)c}{1-b} + \frac{I_1 - ac}{I_2 - c(1-b)} (1-c) \right] I_2$, where $a = \Phi\left(\frac{\eta(q^P)}{\sigma}\right)$, $b = \Phi\left(\frac{\bar{\eta}^P}{\sigma}\right)$, $c = \Phi\left(\frac{\bar{\eta}^P - \Lambda}{\sigma}\right)$, $I_1 = \int_{\bar{\eta}^P+1}^{\eta(q^P)} \Phi\left(\frac{x - \Lambda}{\sigma}\right) dF_x$ and $I_2 = \int_{\bar{\eta}^P}^{+\infty} \Phi\left(\frac{x - \Lambda}{\sigma}\right) dF_x$.

Now we are ready to calculate the PDF by taking a derivative of the CDF:

$$f(\eta^T = x \mid \eta^T > \max\{\bar{\eta}^T, \eta^{-T} + (\bar{\eta}^T - \bar{\eta}^{-T})\}) \\ = \begin{cases} \frac{a'c}{1-b} + \frac{I'_1 - a'c}{I_2 - c(1-b)} (1-c) & \text{if } x > \bar{\eta}^T \\ 0 & \text{otherwise} \end{cases}$$

where $a' = \frac{1}{\sigma}\phi\left(\frac{x}{\sigma}\right)$, $b = \Phi\left(\frac{\bar{\eta}^T}{\sigma}\right)$, $c = \Phi\left(\frac{\bar{\eta}^T - \Lambda}{\sigma}\right)$, $I'_1 = \Phi\left(\frac{x-\Lambda}{\sigma}\right)f_x(x)$ and $I_2 = \int_{\bar{\eta}^T}^{+\infty} \Phi\left(\frac{x-\Lambda}{\sigma}\right)dF_x$. f_x and F_x are the PDF and CDF of $N(0, \sigma^2)$.

Taking into account the fact that book quantities are integers, the ultimate probability to calculate is

$$\Pr(\eta(q^P) \leq \eta^P < \eta(q^P + 1) \mid \eta^P > \max\{\bar{\eta}^P + 1, \eta^E + \Lambda\}) \\ \cdot \Pr(\eta^P > \max\{\bar{\eta}^P + 1, \eta^E + \Lambda\}) \\ = CDF(\eta(q^P + 1)) - CDF(\eta(q^P)) \\ = \left[\frac{\tilde{a}c}{1-b} + \frac{\tilde{I}_1 - \tilde{a}c}{I_2 - c(1-b)} (1-c) \right] I_2$$

where $\tilde{a} = \Phi\left(\frac{\eta(q^P+1)}{\sigma}\right) - \Phi\left(\frac{\eta(q^P)}{\sigma}\right)$ and $\tilde{I}_1 = \int_{\eta(q^P)}^{\eta(q^P+1)} \Phi\left(\frac{x-\Lambda}{\sigma}\right)dF_x$. There are two integrals to calculate: $\tilde{I}_1 = \int_{\eta(q^P)}^{\eta(q^P+1)} \Phi\left(\frac{x-\Lambda}{\sigma}\right)dF_x$ and $I_2 = \int_{\bar{\eta}^P}^{+\infty} \Phi\left(\frac{x-\Lambda}{\sigma}\right)dF_x$. I use Gauss-Chebyshev quadrature with 10 nodes to calculate the first one and Gauss-Laguerre quadrature with 10 nodes to calculate the second one.

Indirect Flow Utility A Kindle nonowner has the following ex-ante indirect flow utility from books:

$$v_i^{book,0} = y_i + \sum_g E\left(\frac{(a_{ig}^P - b_i p_g^P)^2}{2b_i} \mid q_{ig}^P > 0\right) \Pr(q_{ig}^P > 0) \\ = y_i + \frac{1}{2b_i} \sum_g E\left((\eta_{ig}^P - \bar{\eta}_{ig}^P)^2 \mid \eta_{ig}^P - \bar{\eta}_{ig}^P > 0\right) \Pr(\eta_{ig}^P - \bar{\eta}_{ig}^P > 0)$$

where $X \equiv \eta_{ig}^P - \bar{\eta}_{ig}^P \sim N(-\bar{\eta}_{ig}^P, \sigma^2)$ and $\Pr(\eta_{ig}^P - \bar{\eta}_{ig}^P > 0) = 1 - \Phi\left(\frac{\bar{\eta}_{ig}^P}{\sigma}\right)$. From the truncated normal distribution properties, we know that

$$\begin{aligned} E(X^2 | X > 0) &= \text{Var}(X | X > 0) + [E(X | X > 0)]^2 \\ &= \sigma^2 [1 - \lambda(\alpha)(\lambda(\alpha) - \alpha)] + [-\bar{\eta}_{ig}^P + \sigma\lambda(\alpha)]^2 \end{aligned}$$

where $\alpha = \frac{\bar{\eta}_{ig}^P}{\sigma}$ and $\lambda(\alpha) = \frac{\phi(\alpha)}{1 - \Phi(\alpha)}$. This is a closed form solution.

A Kindle owner has the following ex-ante flow utility from books:


$$\begin{aligned} v_i^{book,1} &= y_i + \sum_{T=P,E} \sum_g E\left(\frac{(a_{ig}^T - b_i p_g^T)^2}{2b_i} \mid q_{ig}^T > 0, q_{ig}^{-T} = 0\right) \Pr(q_{ig}^T > 0, q_{ig}^{-T} = 0) \\ &= y_i + \frac{1}{2b_i} \cdot \\ &\quad \sum_g E\left((\eta_{ig}^P - \bar{\eta}_{ig}^P)^2 \mid \eta_{ig}^P > \max\{\bar{\eta}_{ig}^P, \eta_{ig}^E + \Lambda\}\right) \Pr(\eta_{ig}^P > \max\{\bar{\eta}_{ig}^P, \eta_{ig}^E + \Lambda\}) \\ &\quad + E\left((\eta_{ig}^E - \bar{\eta}_{ig}^E)^2 \mid \eta_{ig}^E > \max\{\bar{\eta}_{ig}^E, \eta_{ig}^P - \Lambda\}\right) \Pr(\eta_{ig}^E > \max\{\bar{\eta}_{ig}^E, \eta_{ig}^P - \Lambda\}) \end{aligned}$$

where the probability $\Pr(\eta_{ig}^P > \max\{\bar{\eta}_{ig}^P, \eta_{ig}^E + \Lambda\})$ is already calculated in the last section. To calculate the two conditional expectations, I use the conditional expectation definitions $E[X | \mathcal{H}] = \int_{-\infty}^{+\infty} x \cdot f(x | \mathcal{H}) dx$ and $E[X^2 | \mathcal{H}] = \int_{-\infty}^{+\infty} x^2 \cdot f(x | \mathcal{H}) dx$. The conditional density $f(x | \mathcal{H})$ is calculated in the last section. Given the conditional density, I calculate the conditional expectations using Gauss-Hermite quadrature with 10 nodes. Again, I account for the fact that book quantities are integers when deriving the equations in the final calculation.

C. Computation Algorithm for the Dynamic Pricing Problem

The numerical algorithm is similar to that in Goettler and Gordon (2011). We summarize the algorithm in Figure 10. It contains an inner loop and an outer loop. The inner loop solves the firm and the consumer maximization problem along with the

Figure 10: Computation Algorithm



next period state space given the value function guess. The outer loop updates the value function guess and iterates until convergence.

For each iteration $k = 1, 2, \dots$,

1) Guess the value functions for the firm and the consumers $\{V^{k-1}, W^{k-1}\}$.

2) Given the value function guess, simultaneously solve the firm's first-order conditions at each state. Since the first-order conditions depend on consumers' current choices and next period Δ' , which in turn depend on their rational expectations of Δ' , I solve for a fixed point in Δ' such that consumers' expectations for Δ' are realized according to the state space evolution equation. In particular, to solve for the fixed point, I first guess the next period state space Δ'^{m-1} and the firm's optimal pricing policy $\{P^{m-1}, p^{E,m-1}\}$, where m is the iteration number for the fixed point in the inner loop. Given the guess, I solve the consumers' device adoption problem to get updated next period state space Δ'^m . Given the updated Δ'^m , I solve the firm's first-order conditions at each state and get the updated optimal pricing policy $\{P^m, p^{E,m}\}$. Check convergence of $|\Delta'^m - \Delta'^{m-1}|$, $|P^m - P^{m-1}|$, and $|p^{E,m} - p^{E,m-1}|$. If

converged, let Δ^k and $\{P^k, p^{E,k}\}$ denote this fixed point. This is the solution to the inner loop given the value function guess $\{V^{k-1}, W^{k-1}\}$.

3) Update the value functions given the firm's policy and the next period state space. Denote them $\{V^k, W^k\}$.

4) Check for convergence of the outer loop $|V^k - V^{k-1}|$ and $|W^k - W^{k-1}|$ at the state space grid points Δ . If convergence is not achieved, return to step 2).

Throughout the computation, I discretize the state space evenly into 20 grid points on both dimensions. The range of the state space is between 0 and the initial market size of each type. I use a cubic spline to interpolate between the grid points for the value functions and the policy functions. This is because solving the firm's first-order condition requires differentiable continuation values. The convergence is checked at the grid points.

BIBLIOGRAPHY

- Aguirre, I., Cowan, S., Vickers, J. 2010. "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100(4), 1601-1615(15).
- Besanko, D., Dubé, J. P., Gupta, S. 2003. "Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data," *Management Science*, 49(9), 1121-1138.
- Besanko, D., Winston, W. L. 1990. "Optimal Price Skimming by a Monopolist Facing Rational Consumers," *Management Science*, 36(5), 555-567.
- Chandukala, S. R., Kim, J., Otter, T., Rossi, P. E., Allenby, G. M. 2008. "Choice Models in Marketing: Economic Assumptions, Challenges and Trends," *Foundations and Trends (R) in Marketing*, 2(2), 97-184.
- Chevalier, J. and Goolsbee, A. 2003. "Measuring Prices and Price Competition Online: Amazon and BarnesandNoble.com," *Quantitative Marketing and Economics* 1, 203-222.
- De los Santos, B., Hortaçsu, A. and Wildenbeest, M. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior," *American Economic Review* 102, 2955-2980.
- Derdenger, T., Kumar, V. 2013. "The Dynamic Effects of Bundling as a Product Strategy," *Marketing Science*, 32(6), 827-859.
- Ericson, R., Pakes, A. 1995. "Markov-perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 62(1), 53-82.
- Economides, Seim, Viard. 2008. "Quantifying the Benefits of Entry into Local Phone Service," *RAND Journal of Economics*, RAND Corporation, 39(3), 699-730.
- Gentzkow, M. 2007. "Valuing New Goods in a Model with Complementarity: Online Newspapers," *American Economic Review*, 97(3), 713 - 44.
- Ghose, A., M. Smith, and R. Telang. 2006. "Internet Exchanges for Used Books: An Empirical Analysis of Product Cannibalization and Welfare Impact," *Information Systems Research*, 17(1), 3-19.
- Gil, R., Hartmann, W. R. 2009. "Empirical Analysis of Metering Price Discrimination: Evidence from Concession Sales at Movie Theaters," *Marketing Science*, 28(6), 1046-1062.
- Gilbert, R. 2014. "Ebooks: A Tale of Digital Disruption," working paper. University of California, Berkeley.

- Goettler, R., Gordon, B. 2011, "Does AMD Spur Intel to Innovate More?" *Journal of Political Economy*, 119(6), 1141-1200.
- Gowrisankaran, G., Rysman, M. 2012. "Dynamics of Consumer Demand for New Durable Goods," *Journal of Political Economy*. 120(6), 1173-1219.
- Hartmann, W. R., Nair, H. S. 2010. "Retail Competition and the Dynamics of Demand for Tied Goods," *Marketing Science*, 29(2), 366-386.
- Hendel, I., Nevo, A. 2013. "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, 103(7), 2722-2751.
- Kannan, P. K., Pope, B., Jain, S. 2009. "Pricing Digital Content Product Lines: A Model and Application for the National Academies Press," *Marketing Science*, Lead Article, Vol. 28, No. 4, pp. 620-636.
- Koh, W. T. 2006. "The Micro-foundations of Intertemporal Price Discrimination," *Economic Theory*, 27(2), 393-410.
- Lee, R. 2013. "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," *American Economic Review*, 103(7), 2960-3000.
- Lazarev, J. 2013. "The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in U.S. Monopoly Markets," working paper.
- Leung, H. M. 1997. "Intertemporal Price Discrimination and Consumer Demand," *Journal of Economics*, 65(1), 19-40.
- Liu, H. 2010. "Dynamics of Pricing in the Video Game Console Market: Skimming or Penetration?" *Journal of Marketing Research*, 47(3), 428-443.
- Melnikov. 2013. "Demand for Differentiated Durable Products: The Case of the US Computer Printer Market," *Economic Inquiry*, Vol. 51.
- Nair, H. 2007. "Intertemporal Price Discrimination with Forward-looking Consumers: Application to the US Market for Console Video-games," *Quantitative Marketing and Economics*, 5(3), 239-292.
- Oberholzer-Gee, F., Strumpf, K. 2007. "The Effect of File Sharing on Record Sales: An Empirical Analysis," *Journal of Political Economy*, Vol. 115, No. 1, pp.1-42.
- Reimers, I. and Waldfogel, J. 2014. "Throwing the Books at Them: Amazon's Puzzling Long Run Pricing Strategy," Mimeo.
- Rust, J. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999-1033.

- Sriram, S., Chintagunta, P. K., Agarwal, M. K. 2010. "Investigating Consumer Purchase Behavior in Related Technology Product Categories," *Marketing Science*, 29(2), 291-314.
- Stokey, N. L. 1979. "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, 93(3), 355-371.
- Stokey, N. L. 1981, "Rational Expectations and Durable Goods Pricing," *The Bell Journal of Economics* 12(1), 112-128.
- Waldfoegel, J. 2007. "'Lost' on the Web: Does Web Distribution Stimulate or Depress Television Viewing?" *NBER Working Papers* 13497, National Bureau of Economic Research, Inc.